

US008669939B2

(12) **United States Patent**
Underkoffler et al.

(10) **Patent No.:** **US 8,669,939 B2**
(45) **Date of Patent:** **Mar. 11, 2014**

(54) **SPATIAL, MULTI-MODAL CONTROL DEVICE FOR USE WITH SPATIAL OPERATING SYSTEM**

(75) Inventors: **John S. Underkoffler**, Los Angeles, CA (US); **Carlton Sparrell**, Marblehead, MA (US); **Harald Belker**, Los Angeles, CA (US); **Kwindla Hultman Kramer**, Los Angeles, CA (US)

(73) Assignee: **Oblong Industries, Inc.**, Los Angeles, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 895 days.

(21) Appl. No.: **12/789,262**

(22) Filed: **May 27, 2010**

(65) **Prior Publication Data**

US 2011/0025598 A1 Feb. 3, 2011

Related U.S. Application Data

(63) Continuation-in-part of application No. 11/350,697, filed on Feb. 8, 2006, now Pat. No. 7,598,942, and a continuation-in-part of application No. 12/109,263, filed on Apr. 24, 2008, now Pat. No. 8,407,725, and a continuation-in-part of application No. 12/553,845, filed on Sep. 3, 2009, now Pat. No. 8,531,396.

(60) Provisional application No. 61/181,621, filed on May 27, 2009.

(51) **Int. Cl.**
G09G 5/08 (2006.01)

(52) **U.S. Cl.**
USPC **345/158**; 345/169

(58) **Field of Classification Search**
USPC 345/156-169, 173-179, 204-215; 178/18.01-18.09, 19.01-19.07, 178/20.01-20.04

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,229,017 B2 *	6/2007	Richley et al.	235/385
8,212,550 B2 *	7/2012	Katsurahira et al.	324/207.11
8,363,098 B2 *	1/2013	Rosener et al.	348/77

* cited by examiner

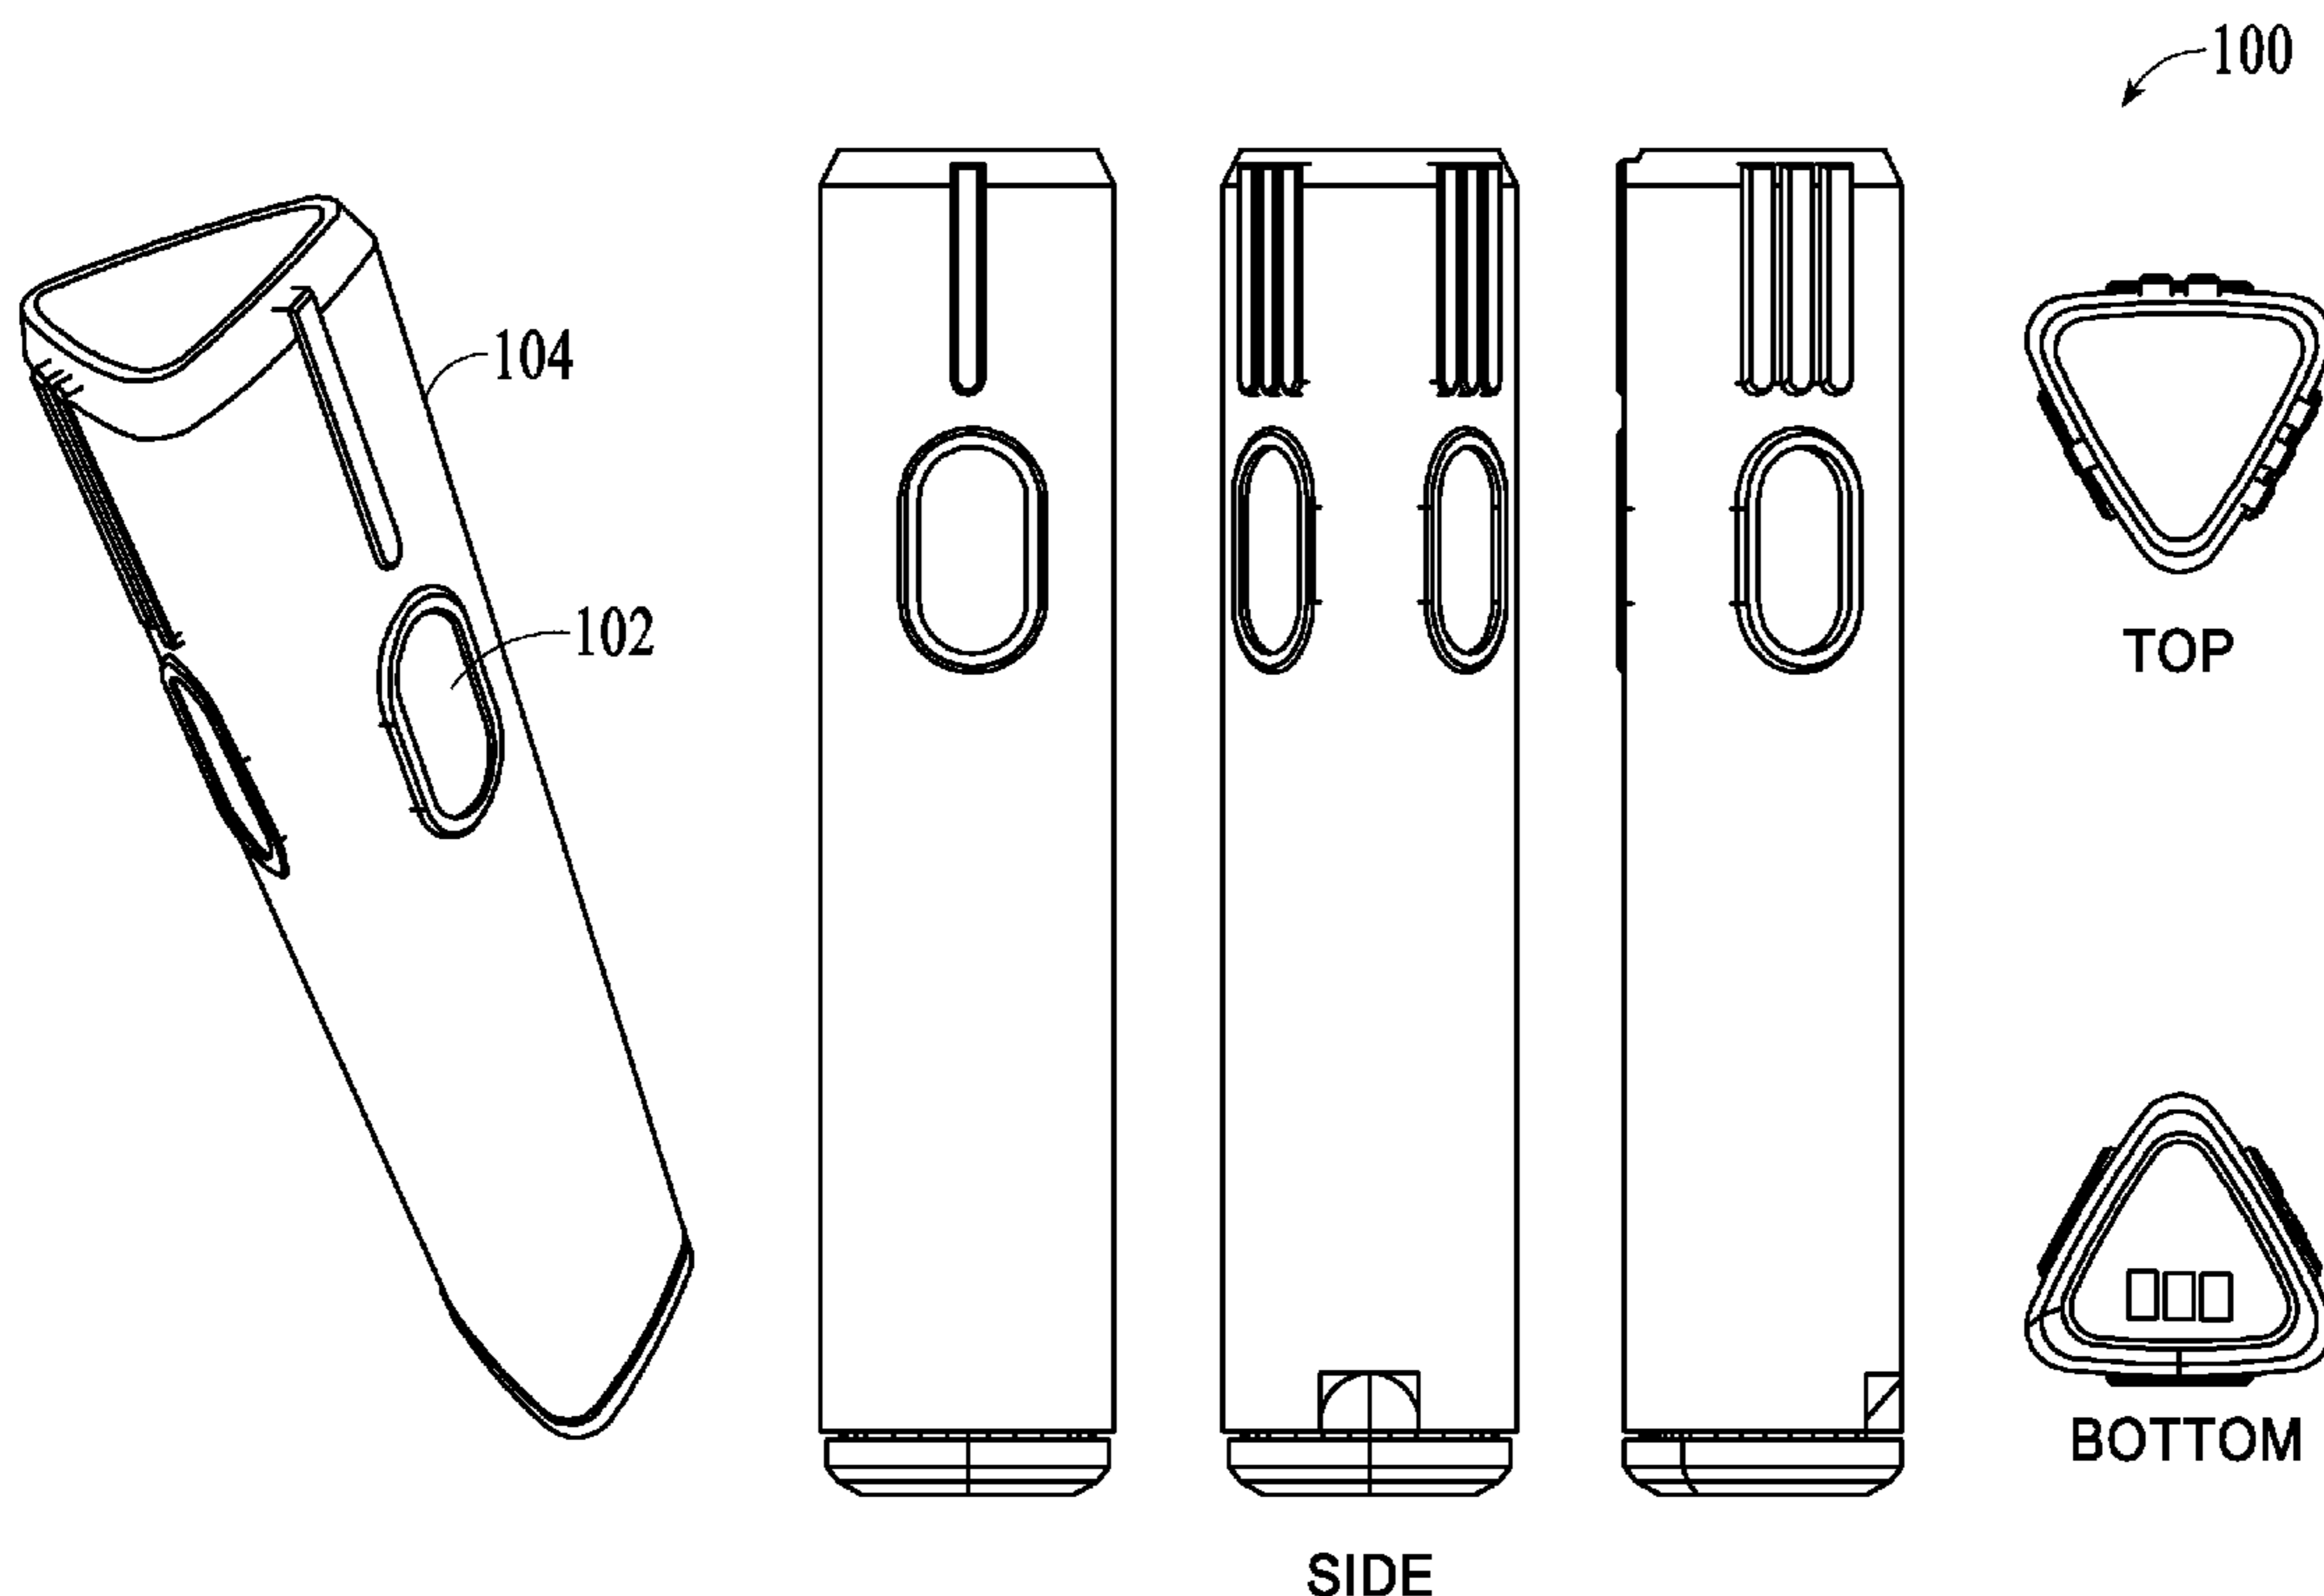
Primary Examiner — Vijay Shankar

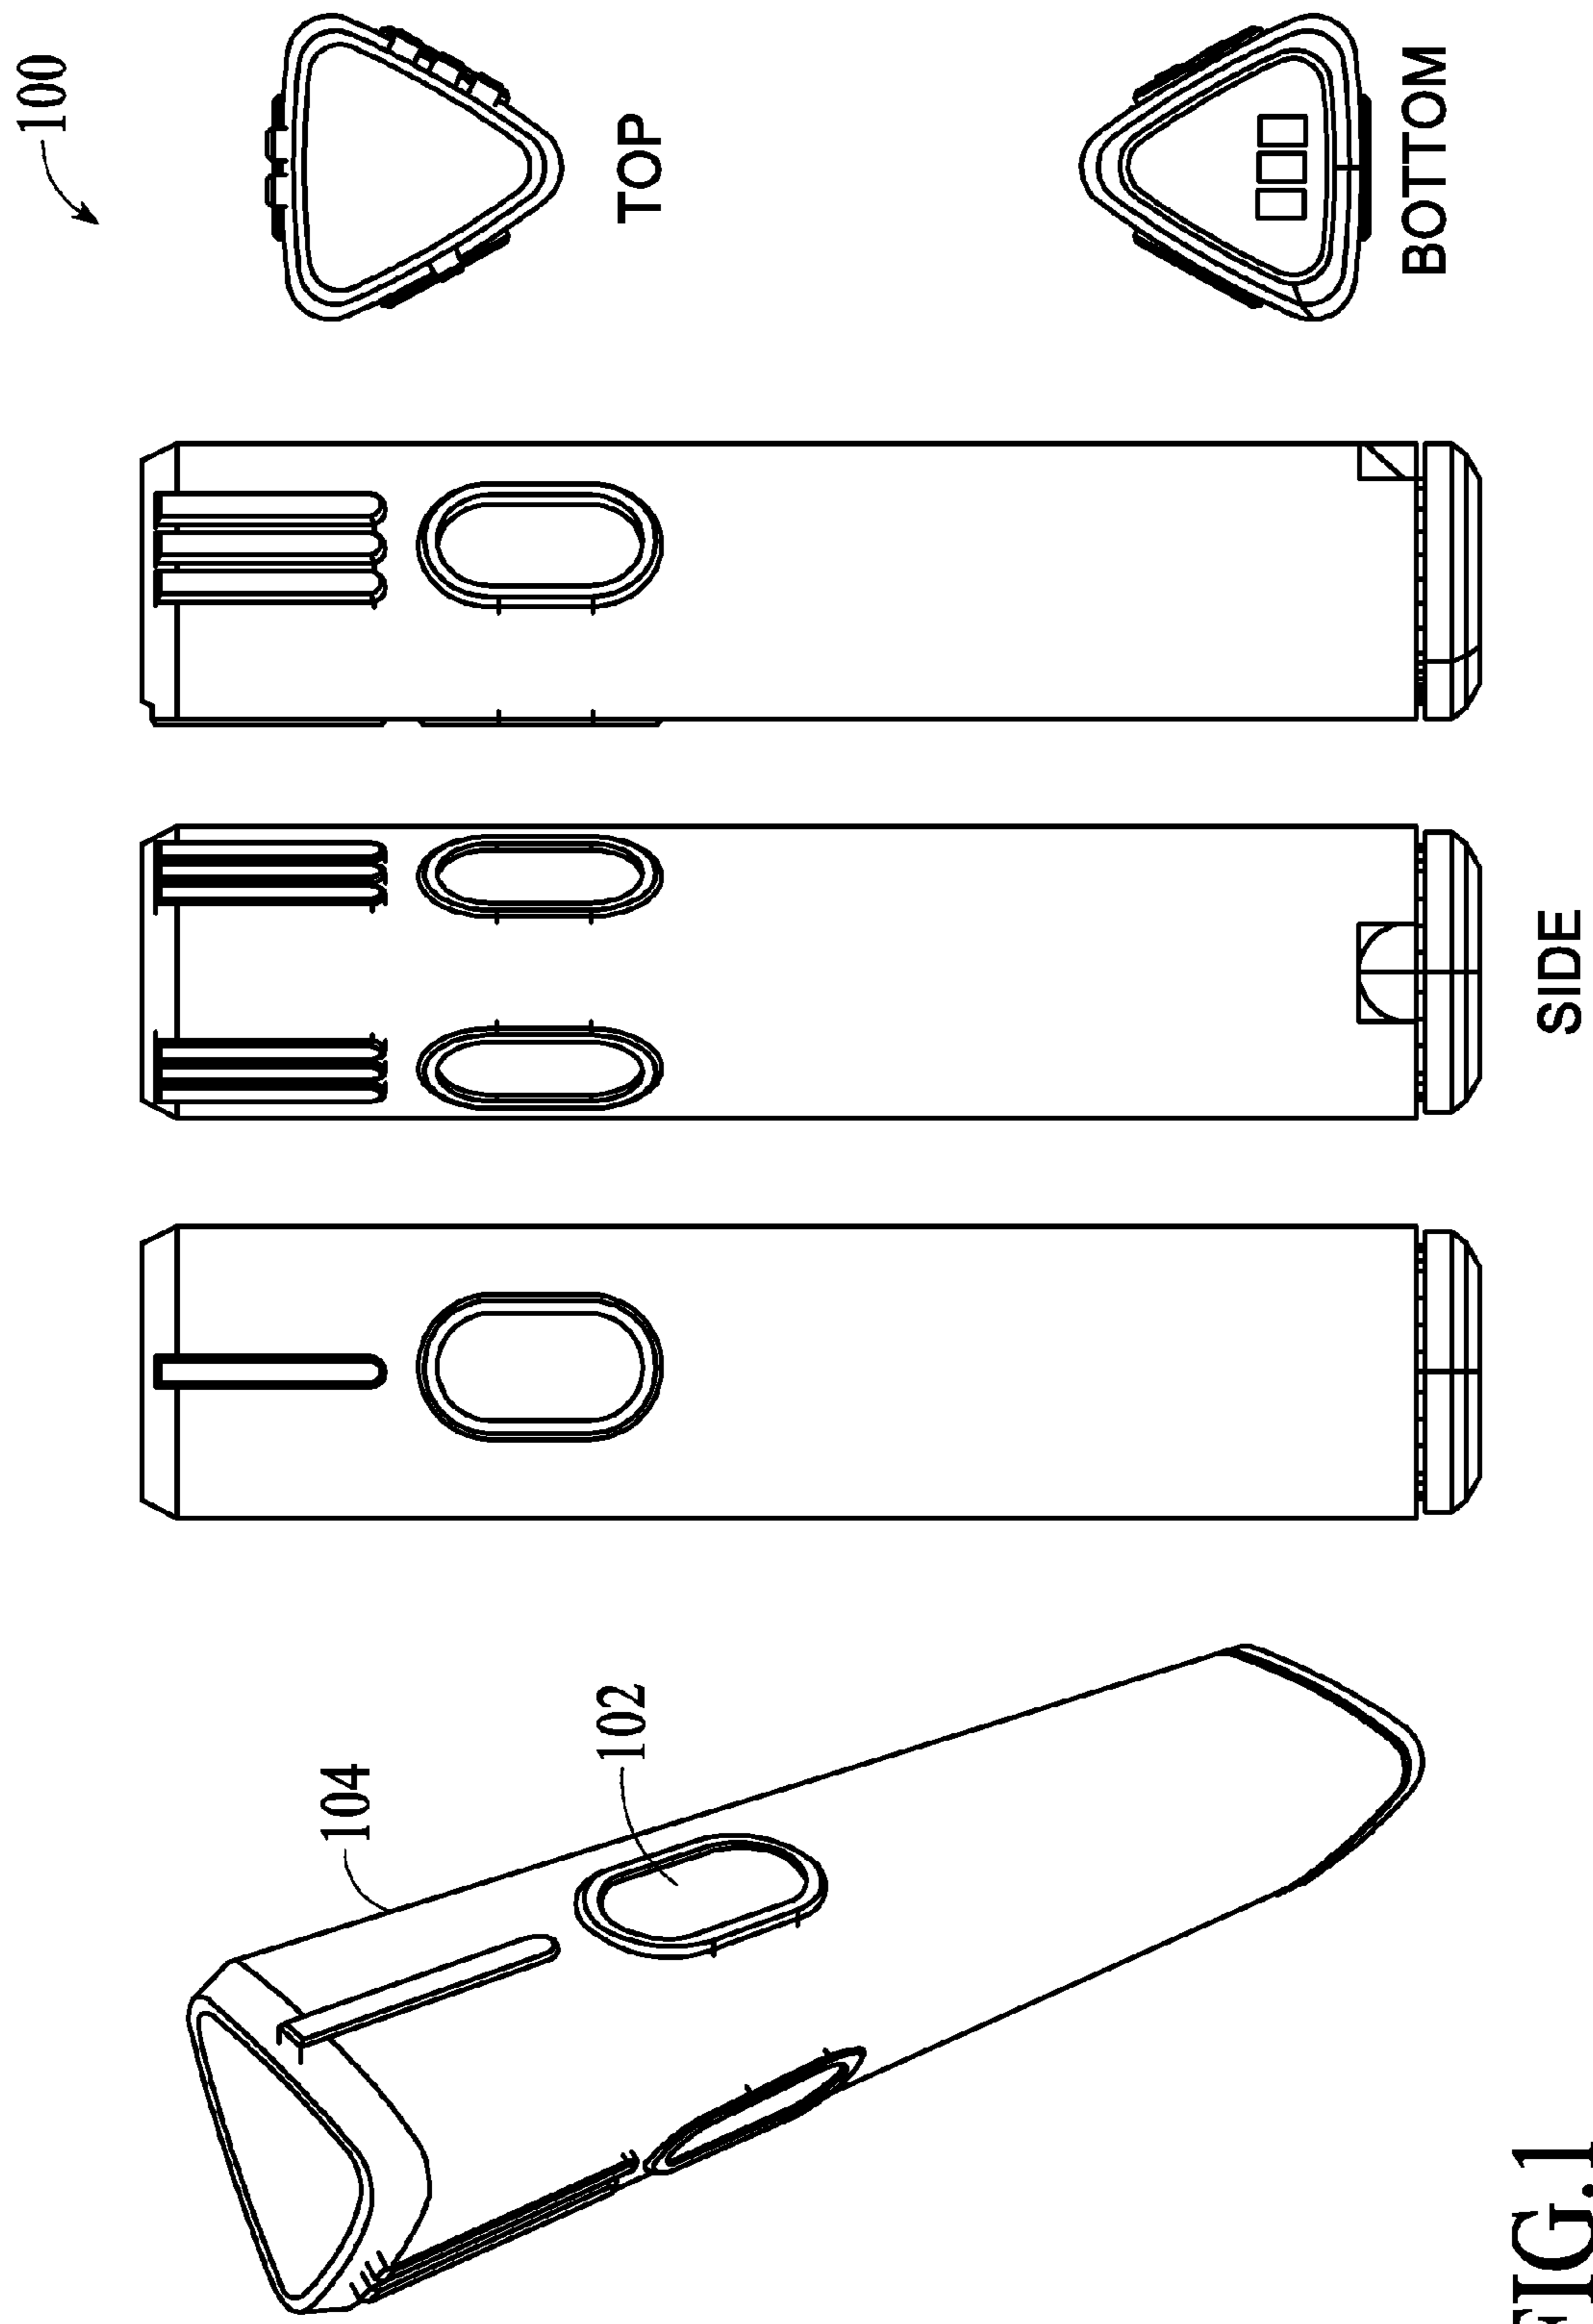
(74) *Attorney, Agent, or Firm* — Gregory & Sawrie LLP

(57) **ABSTRACT**

A system comprising an input device includes a detector coupled to a processor. The detector detects an orientation of the input device. The input device has multiple modal orientations corresponding to the orientation. The modal orientations correspond to multiple input modes of a gestural control system. The detector is coupled to the gestural control system and automatically controls selection of an input mode in response to the orientation.

57 Claims, 26 Drawing Sheets





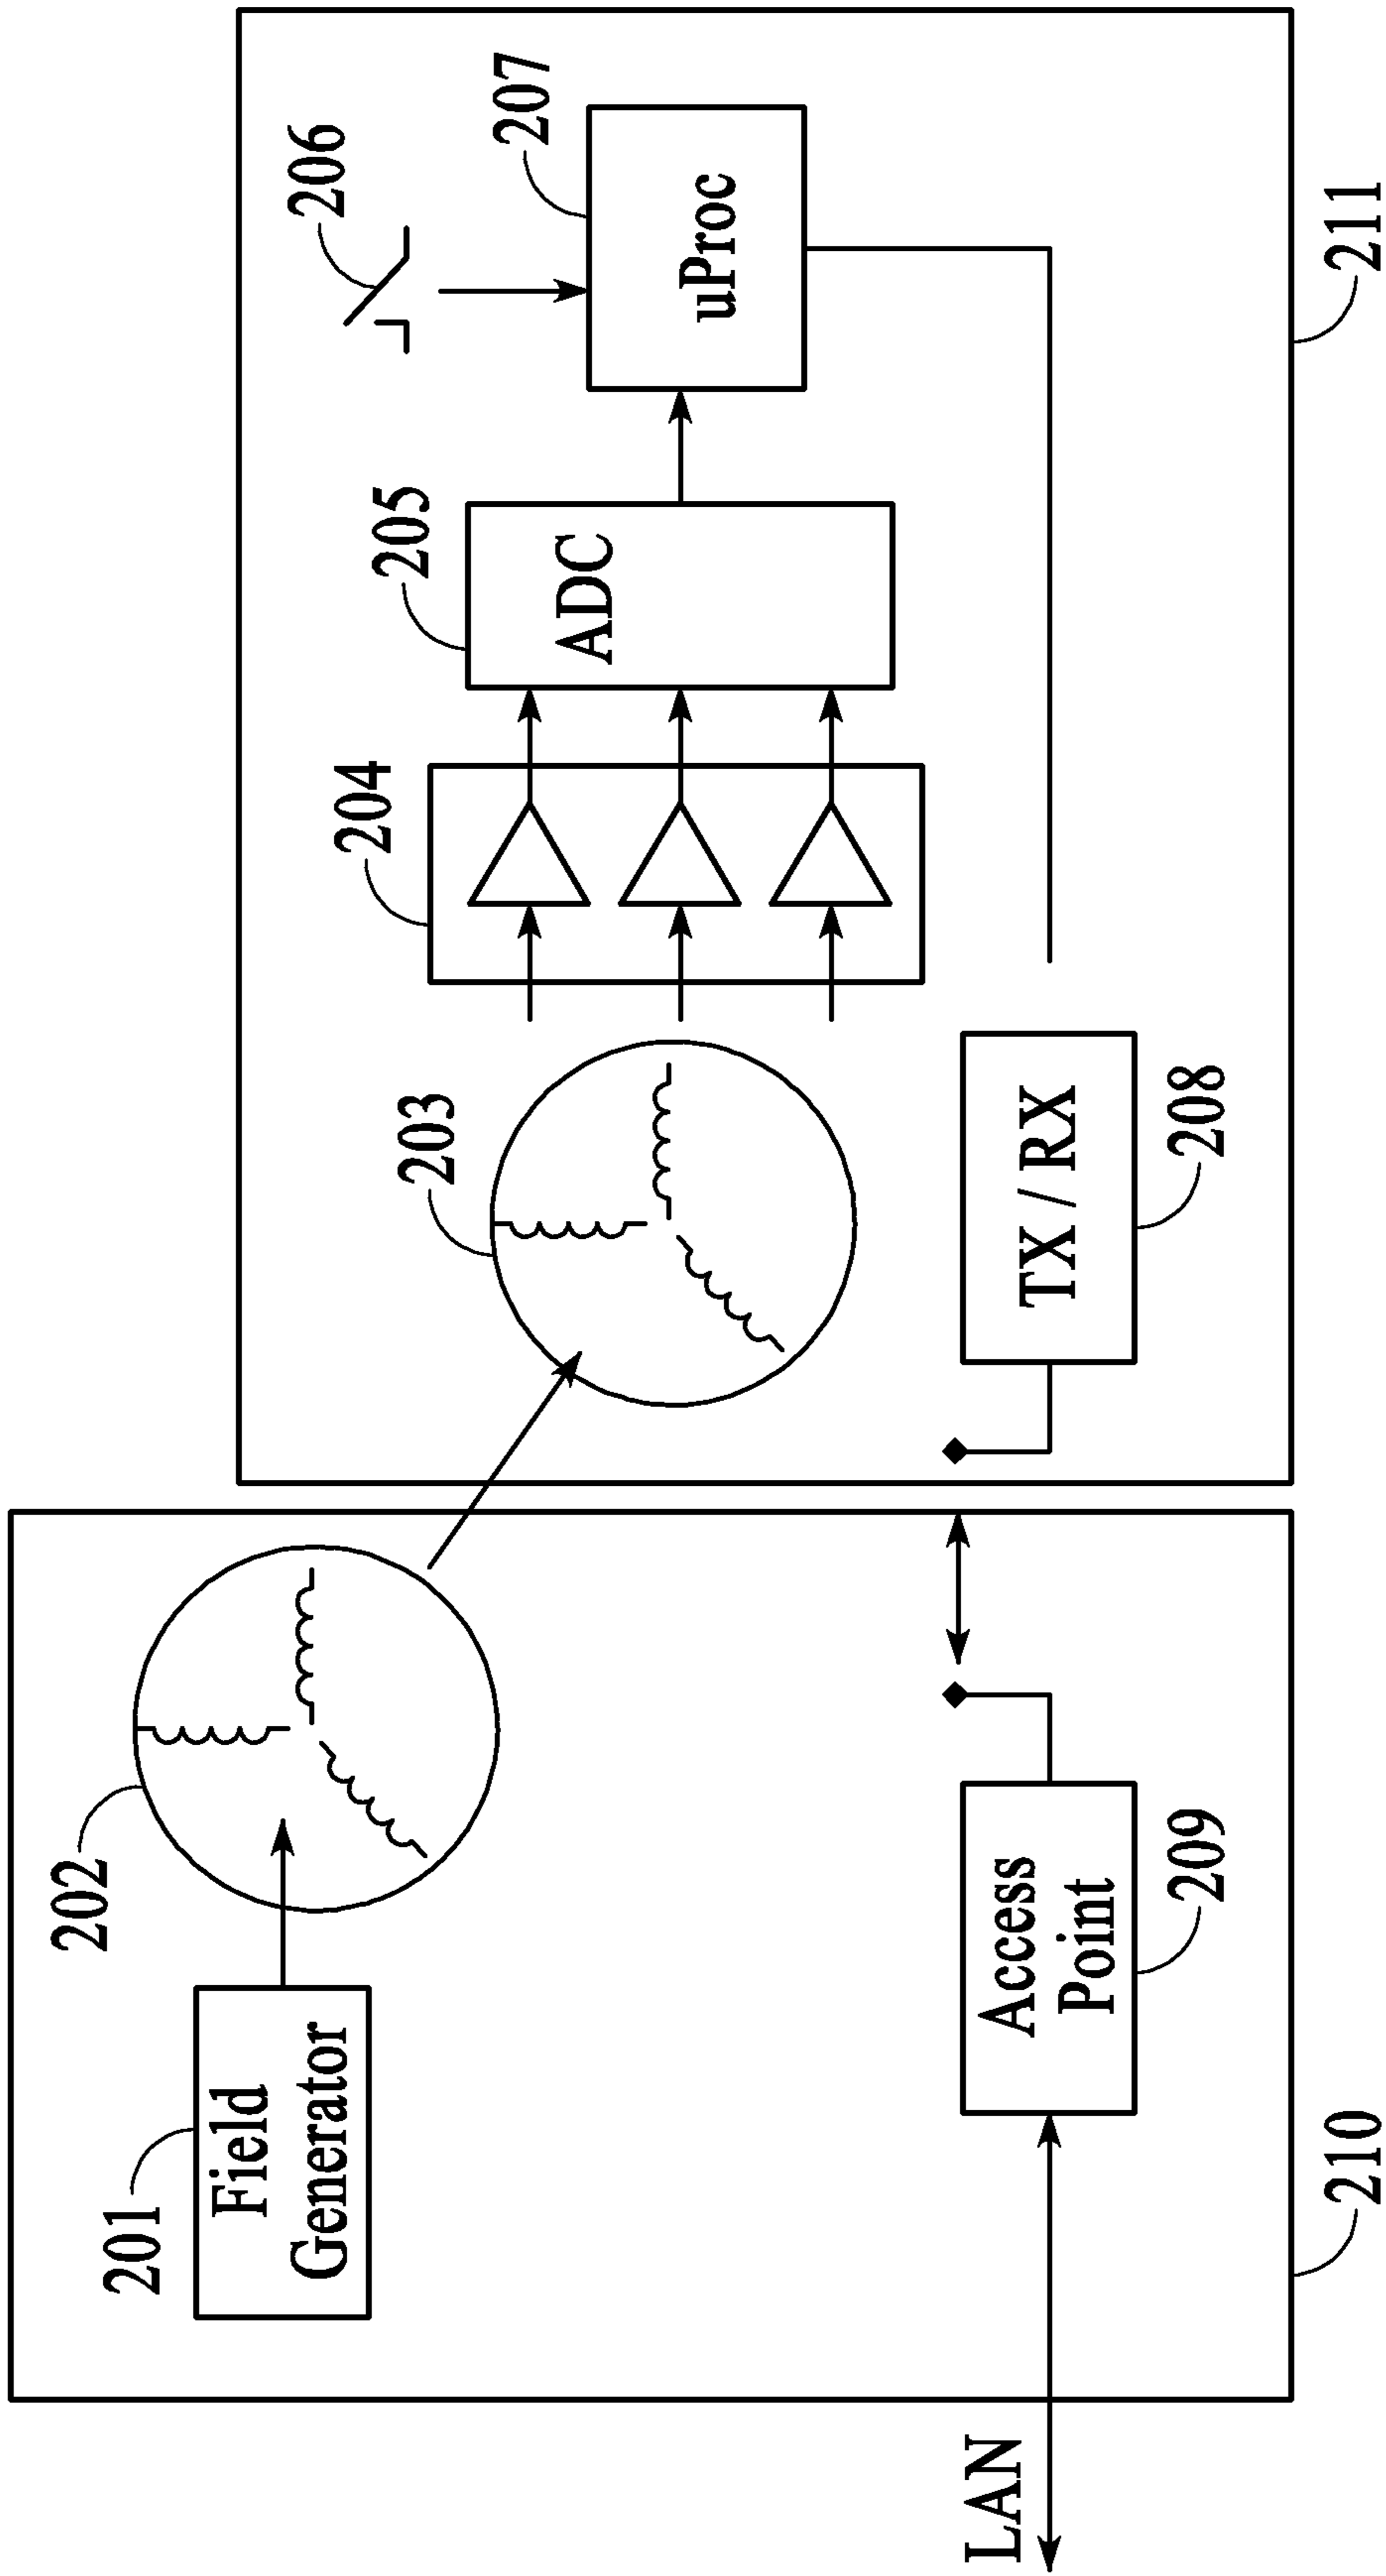


FIG.2

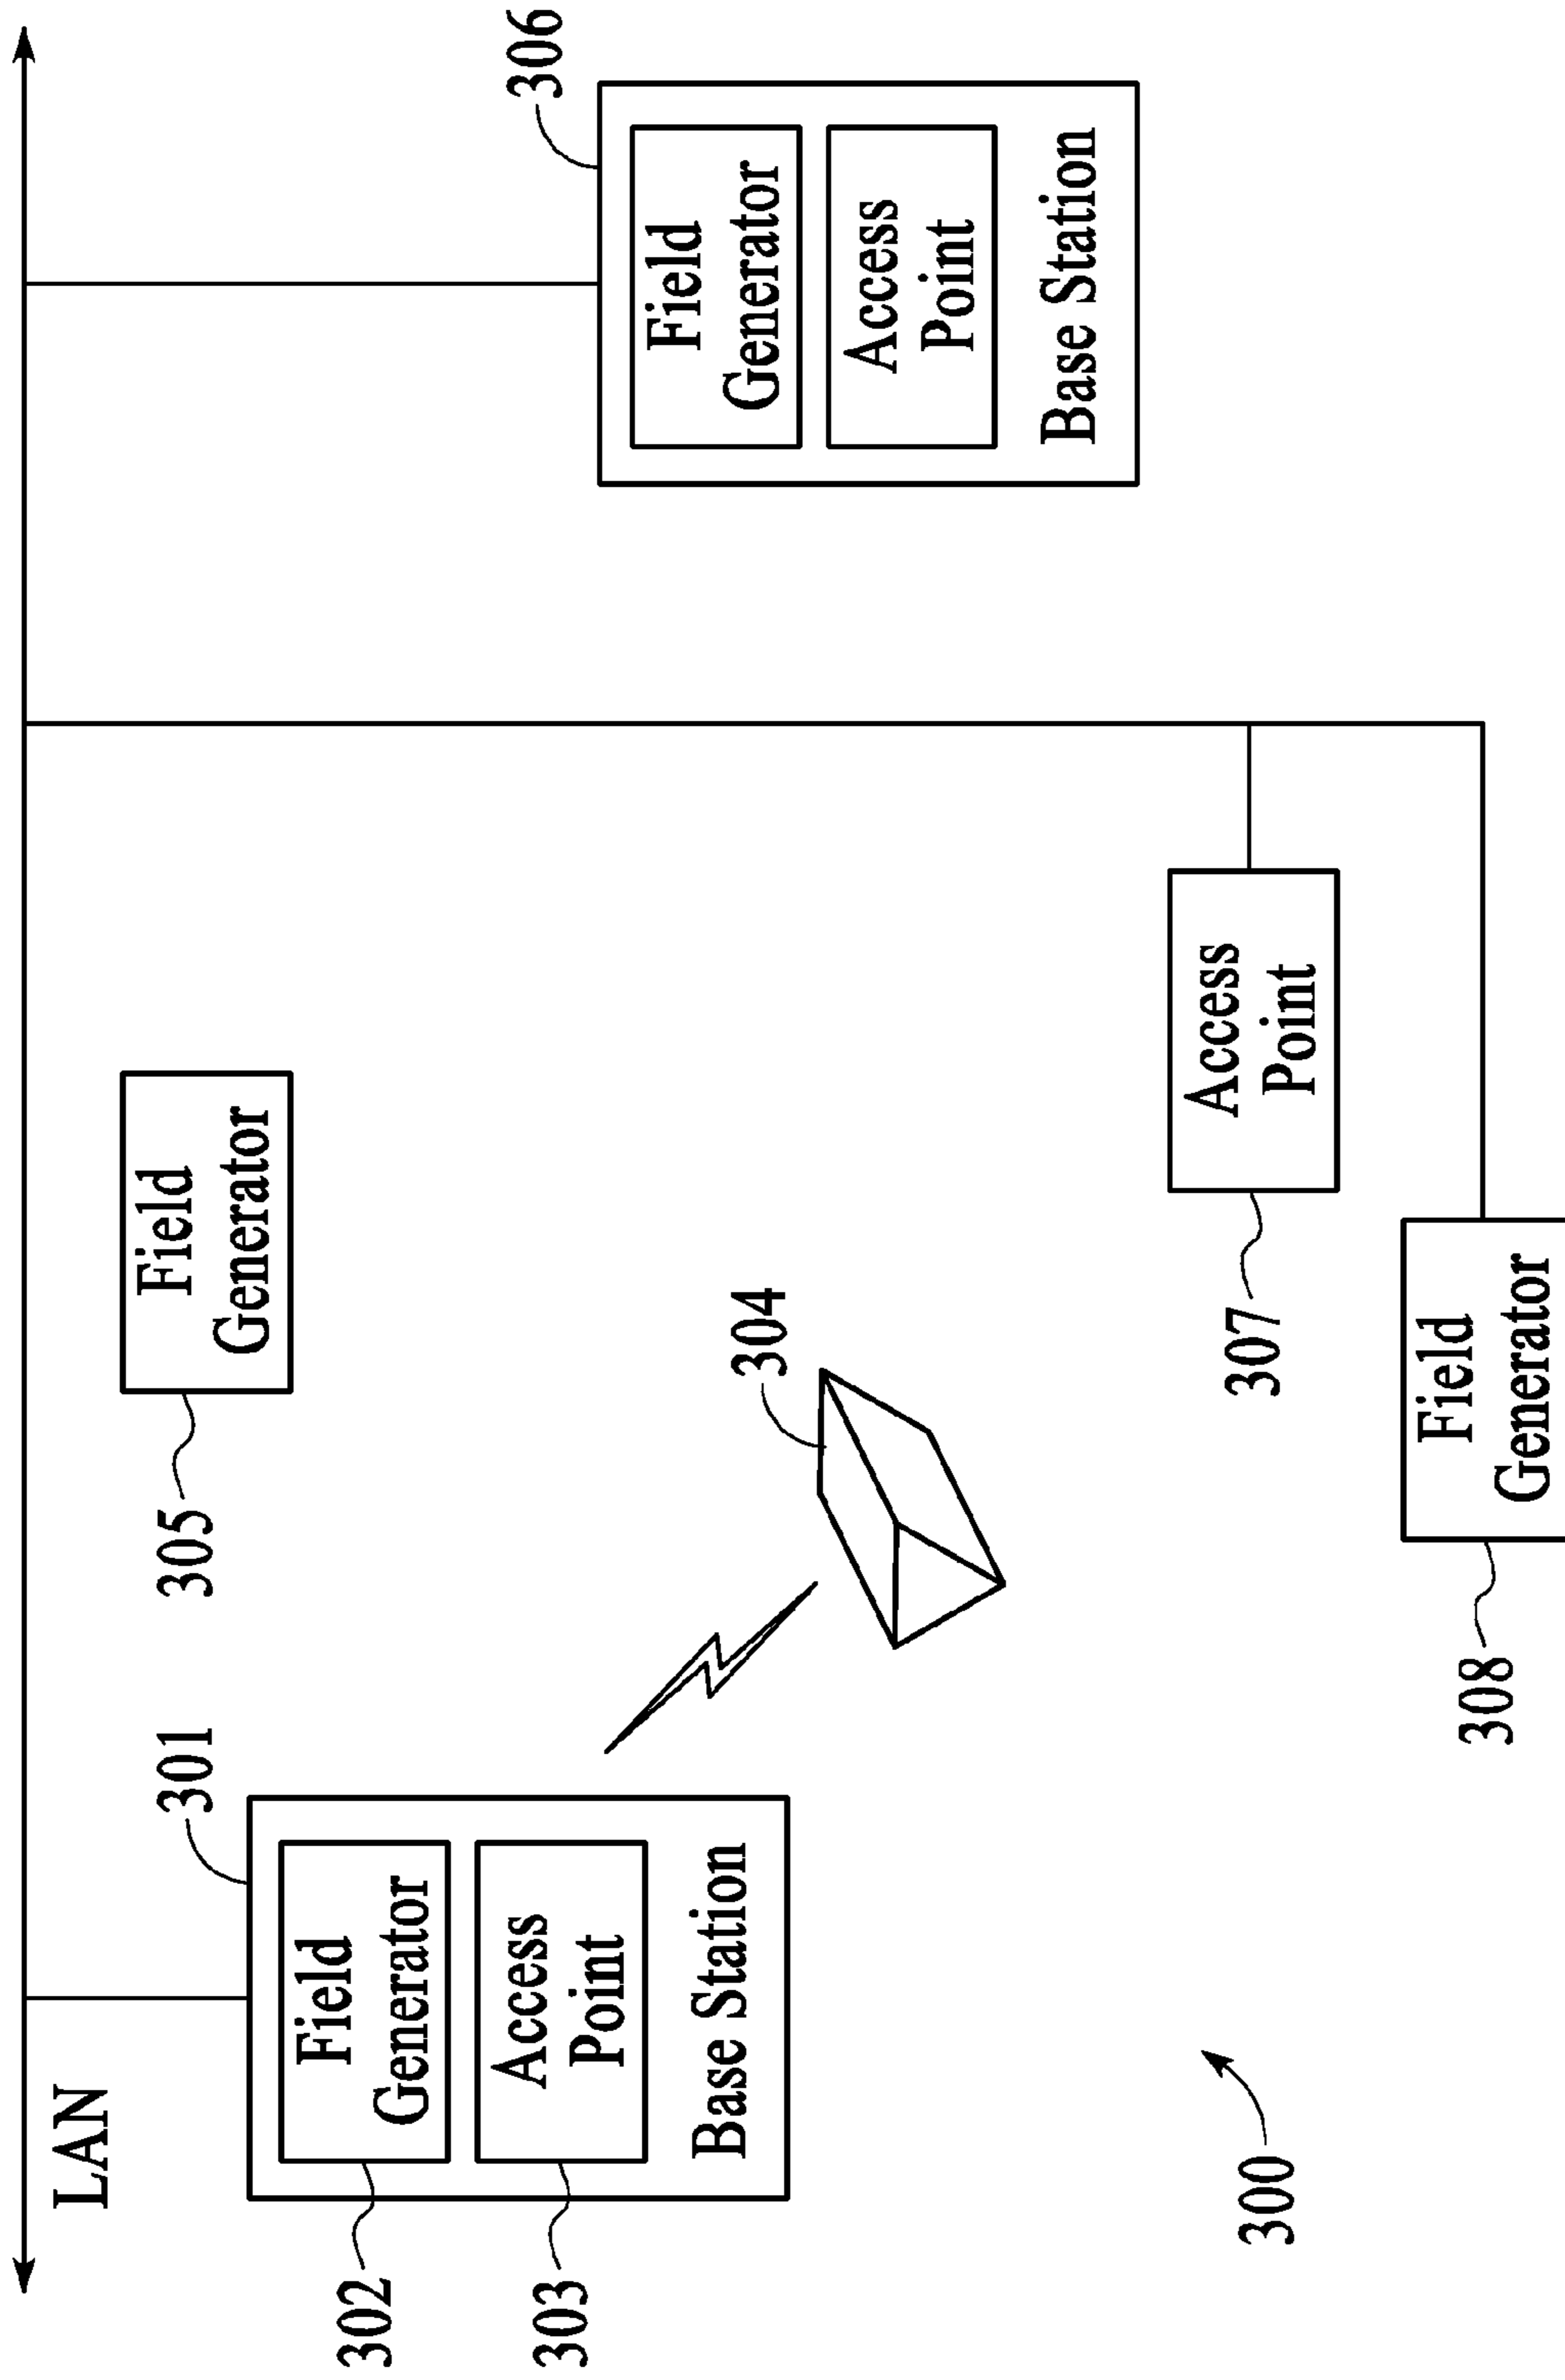


FIG.3

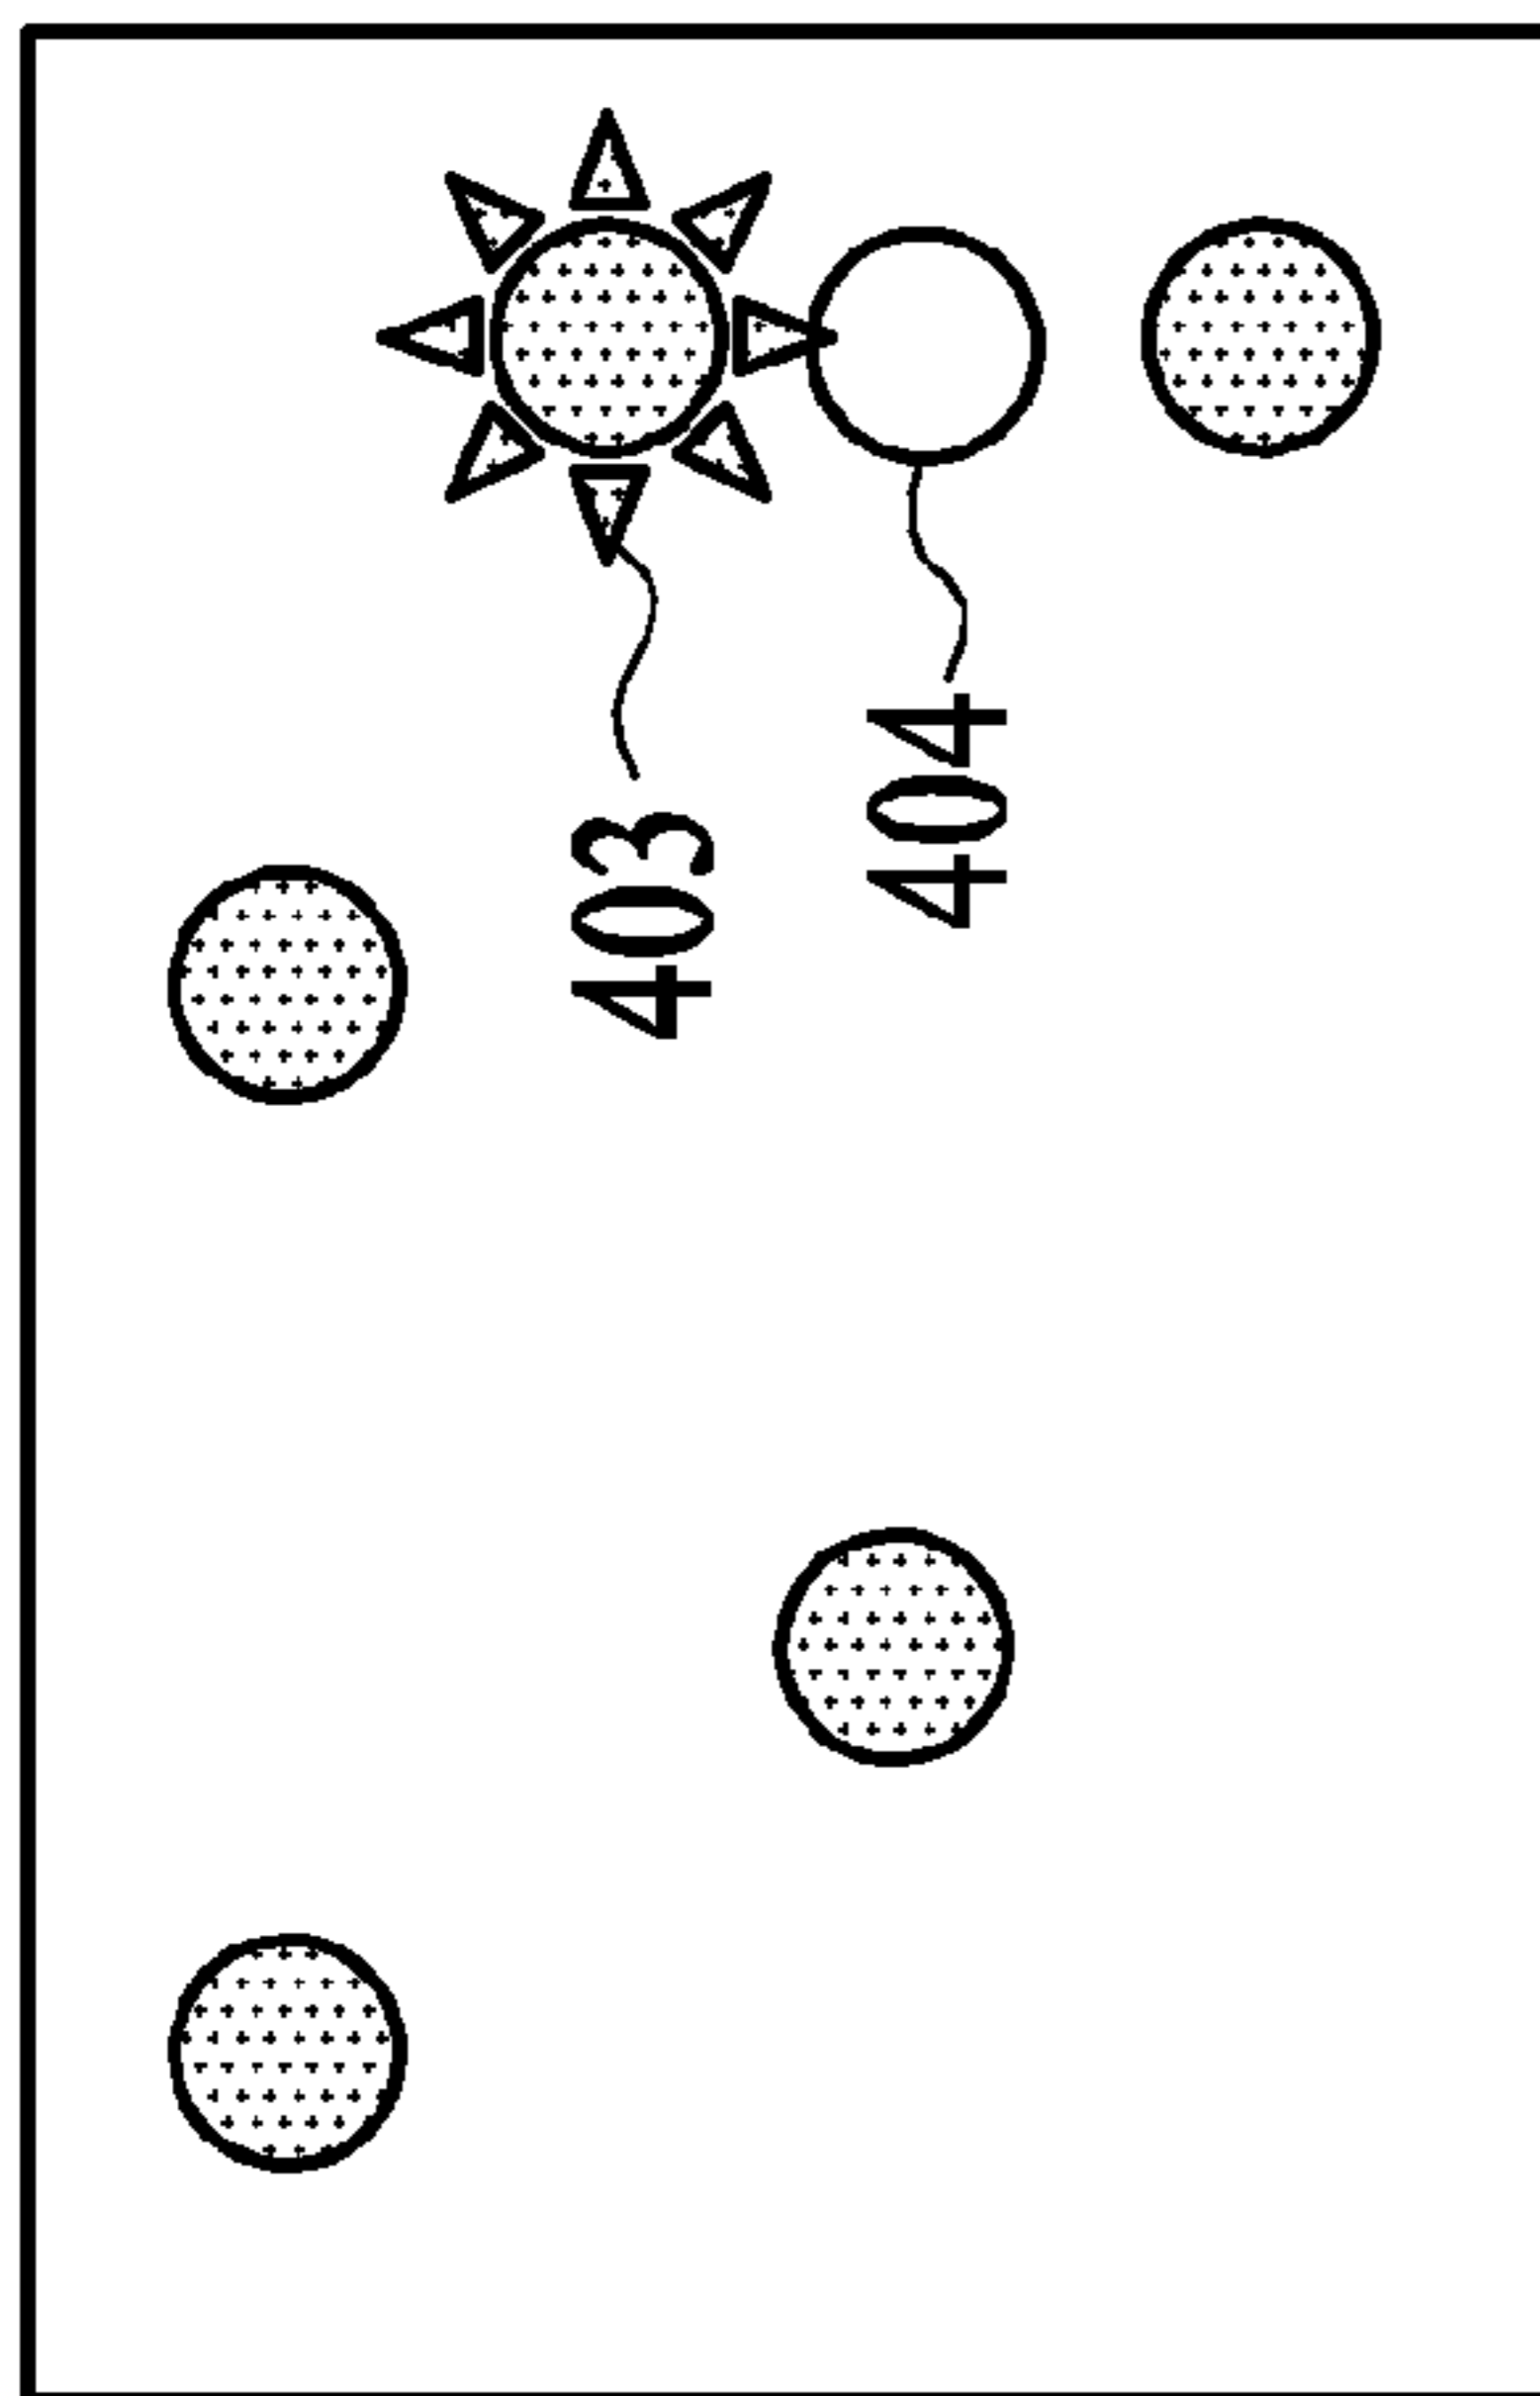


FIG. 4B

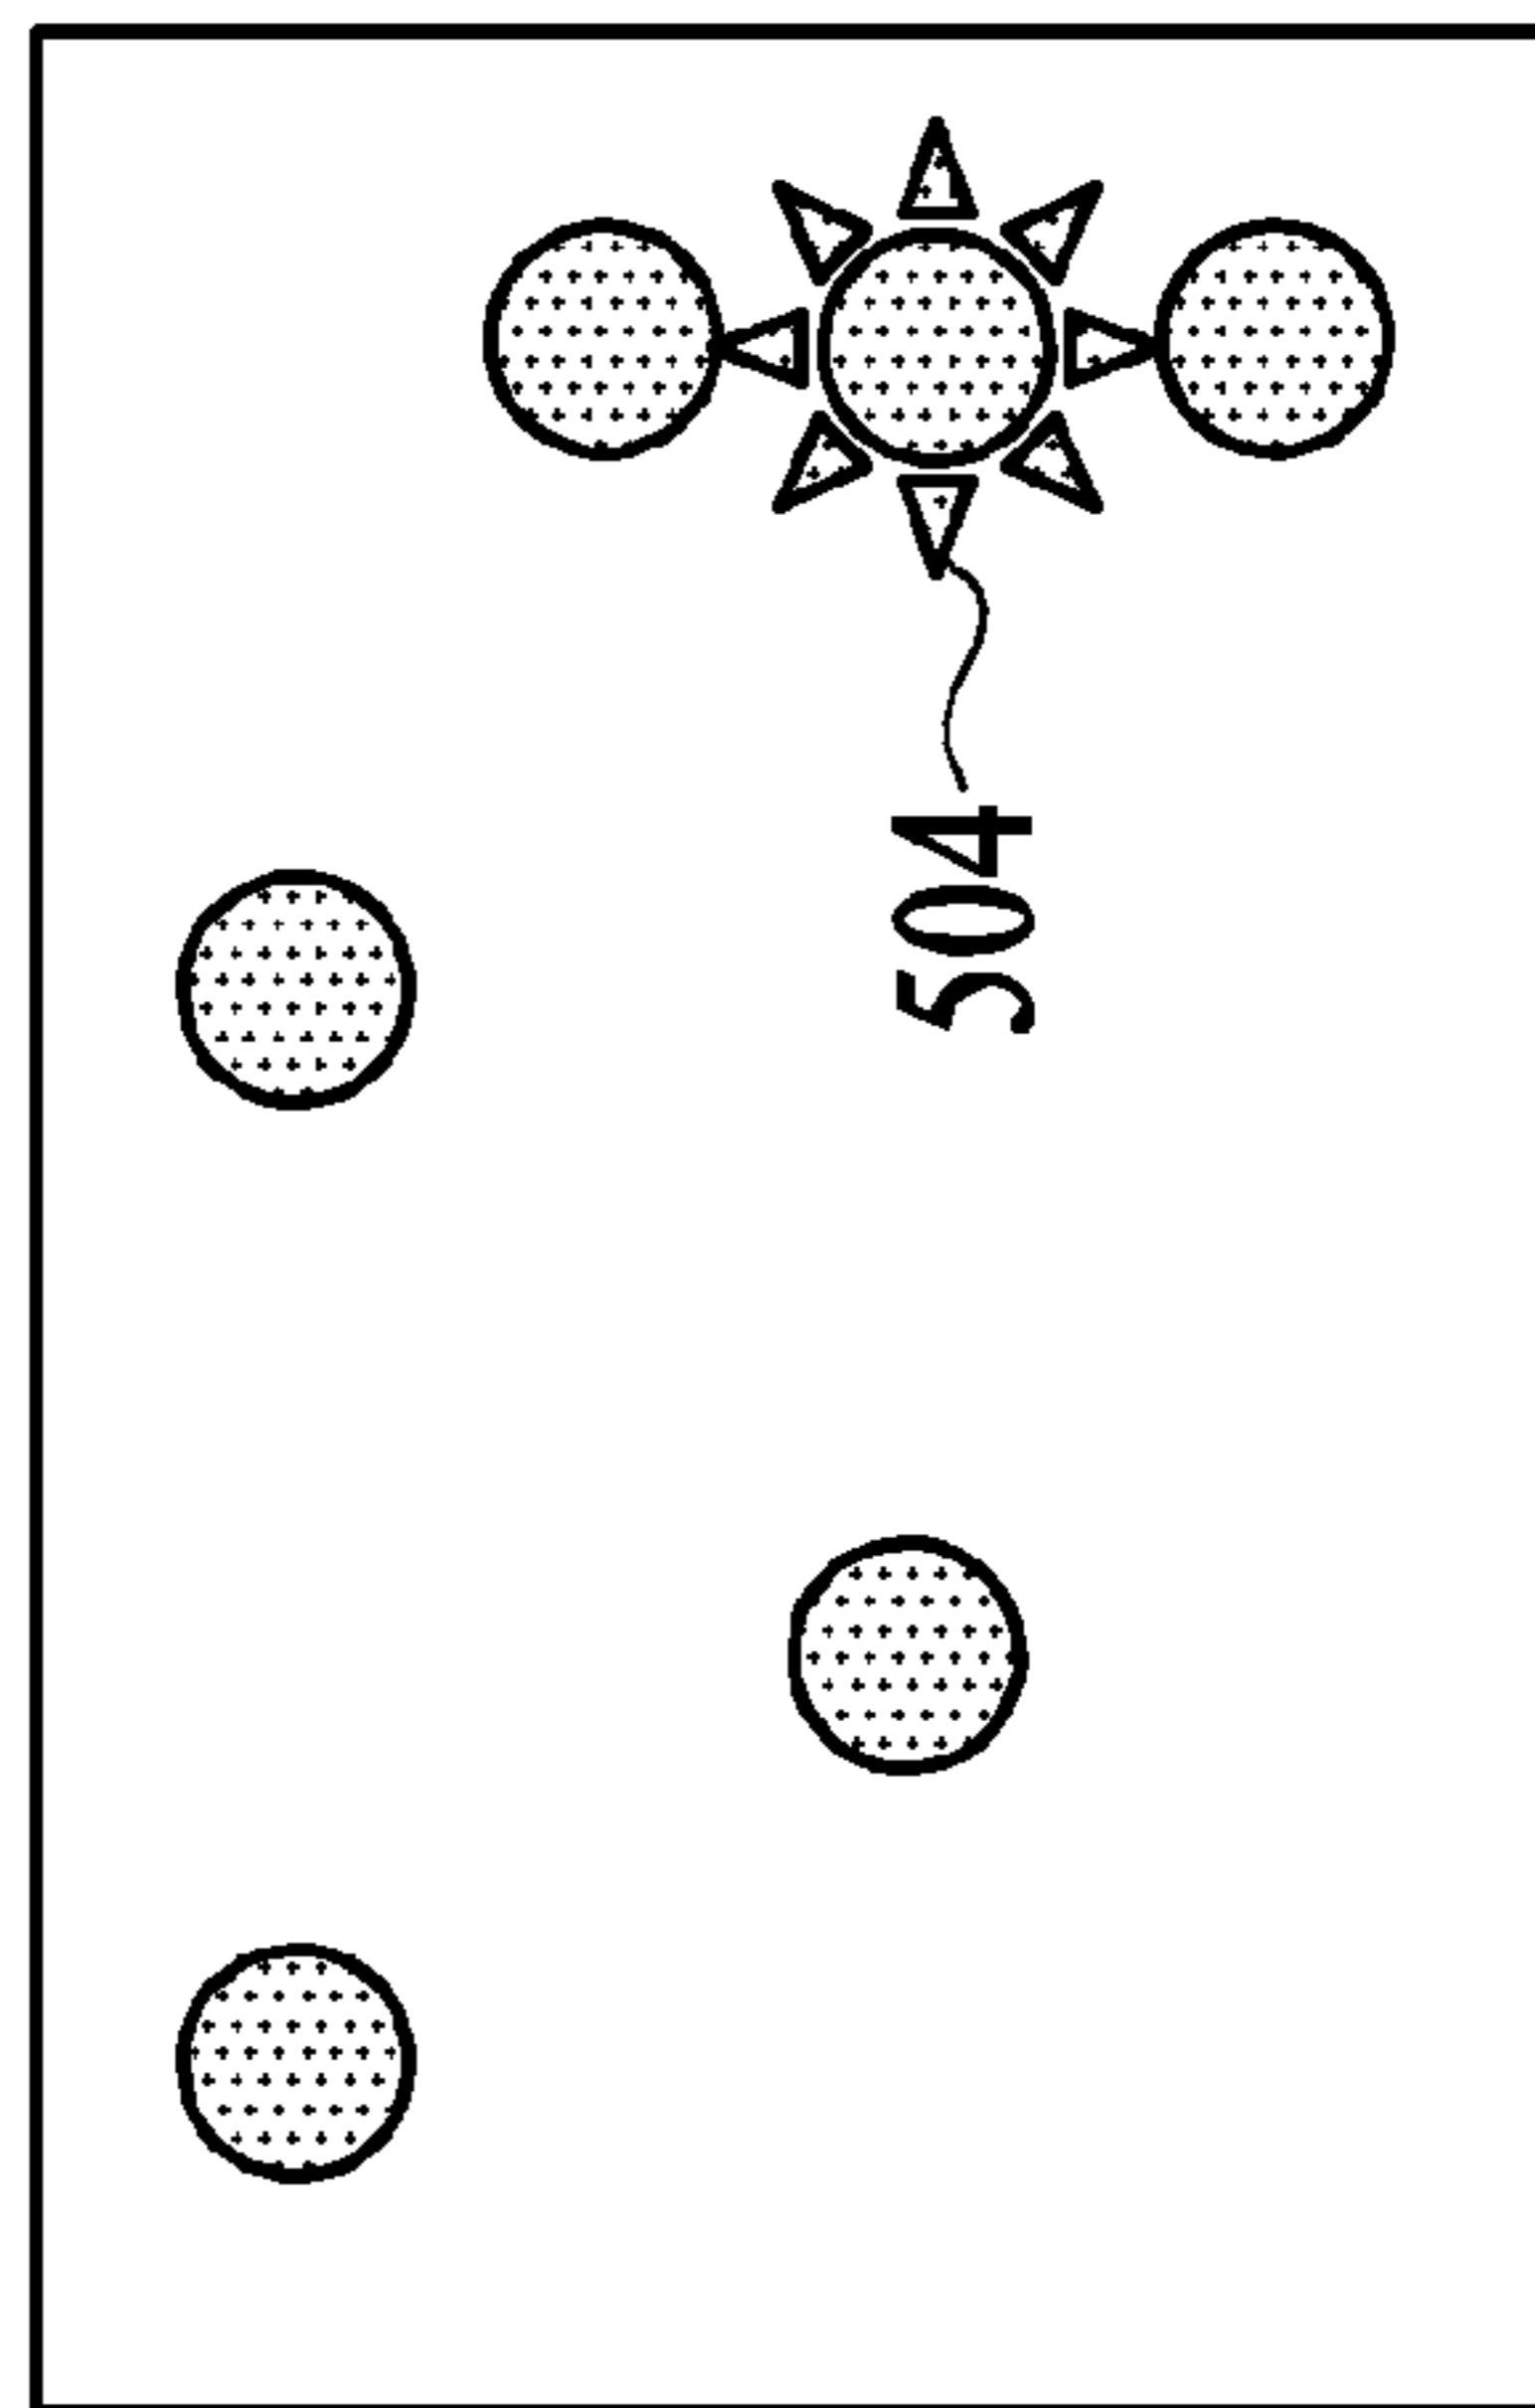


FIG. 5B

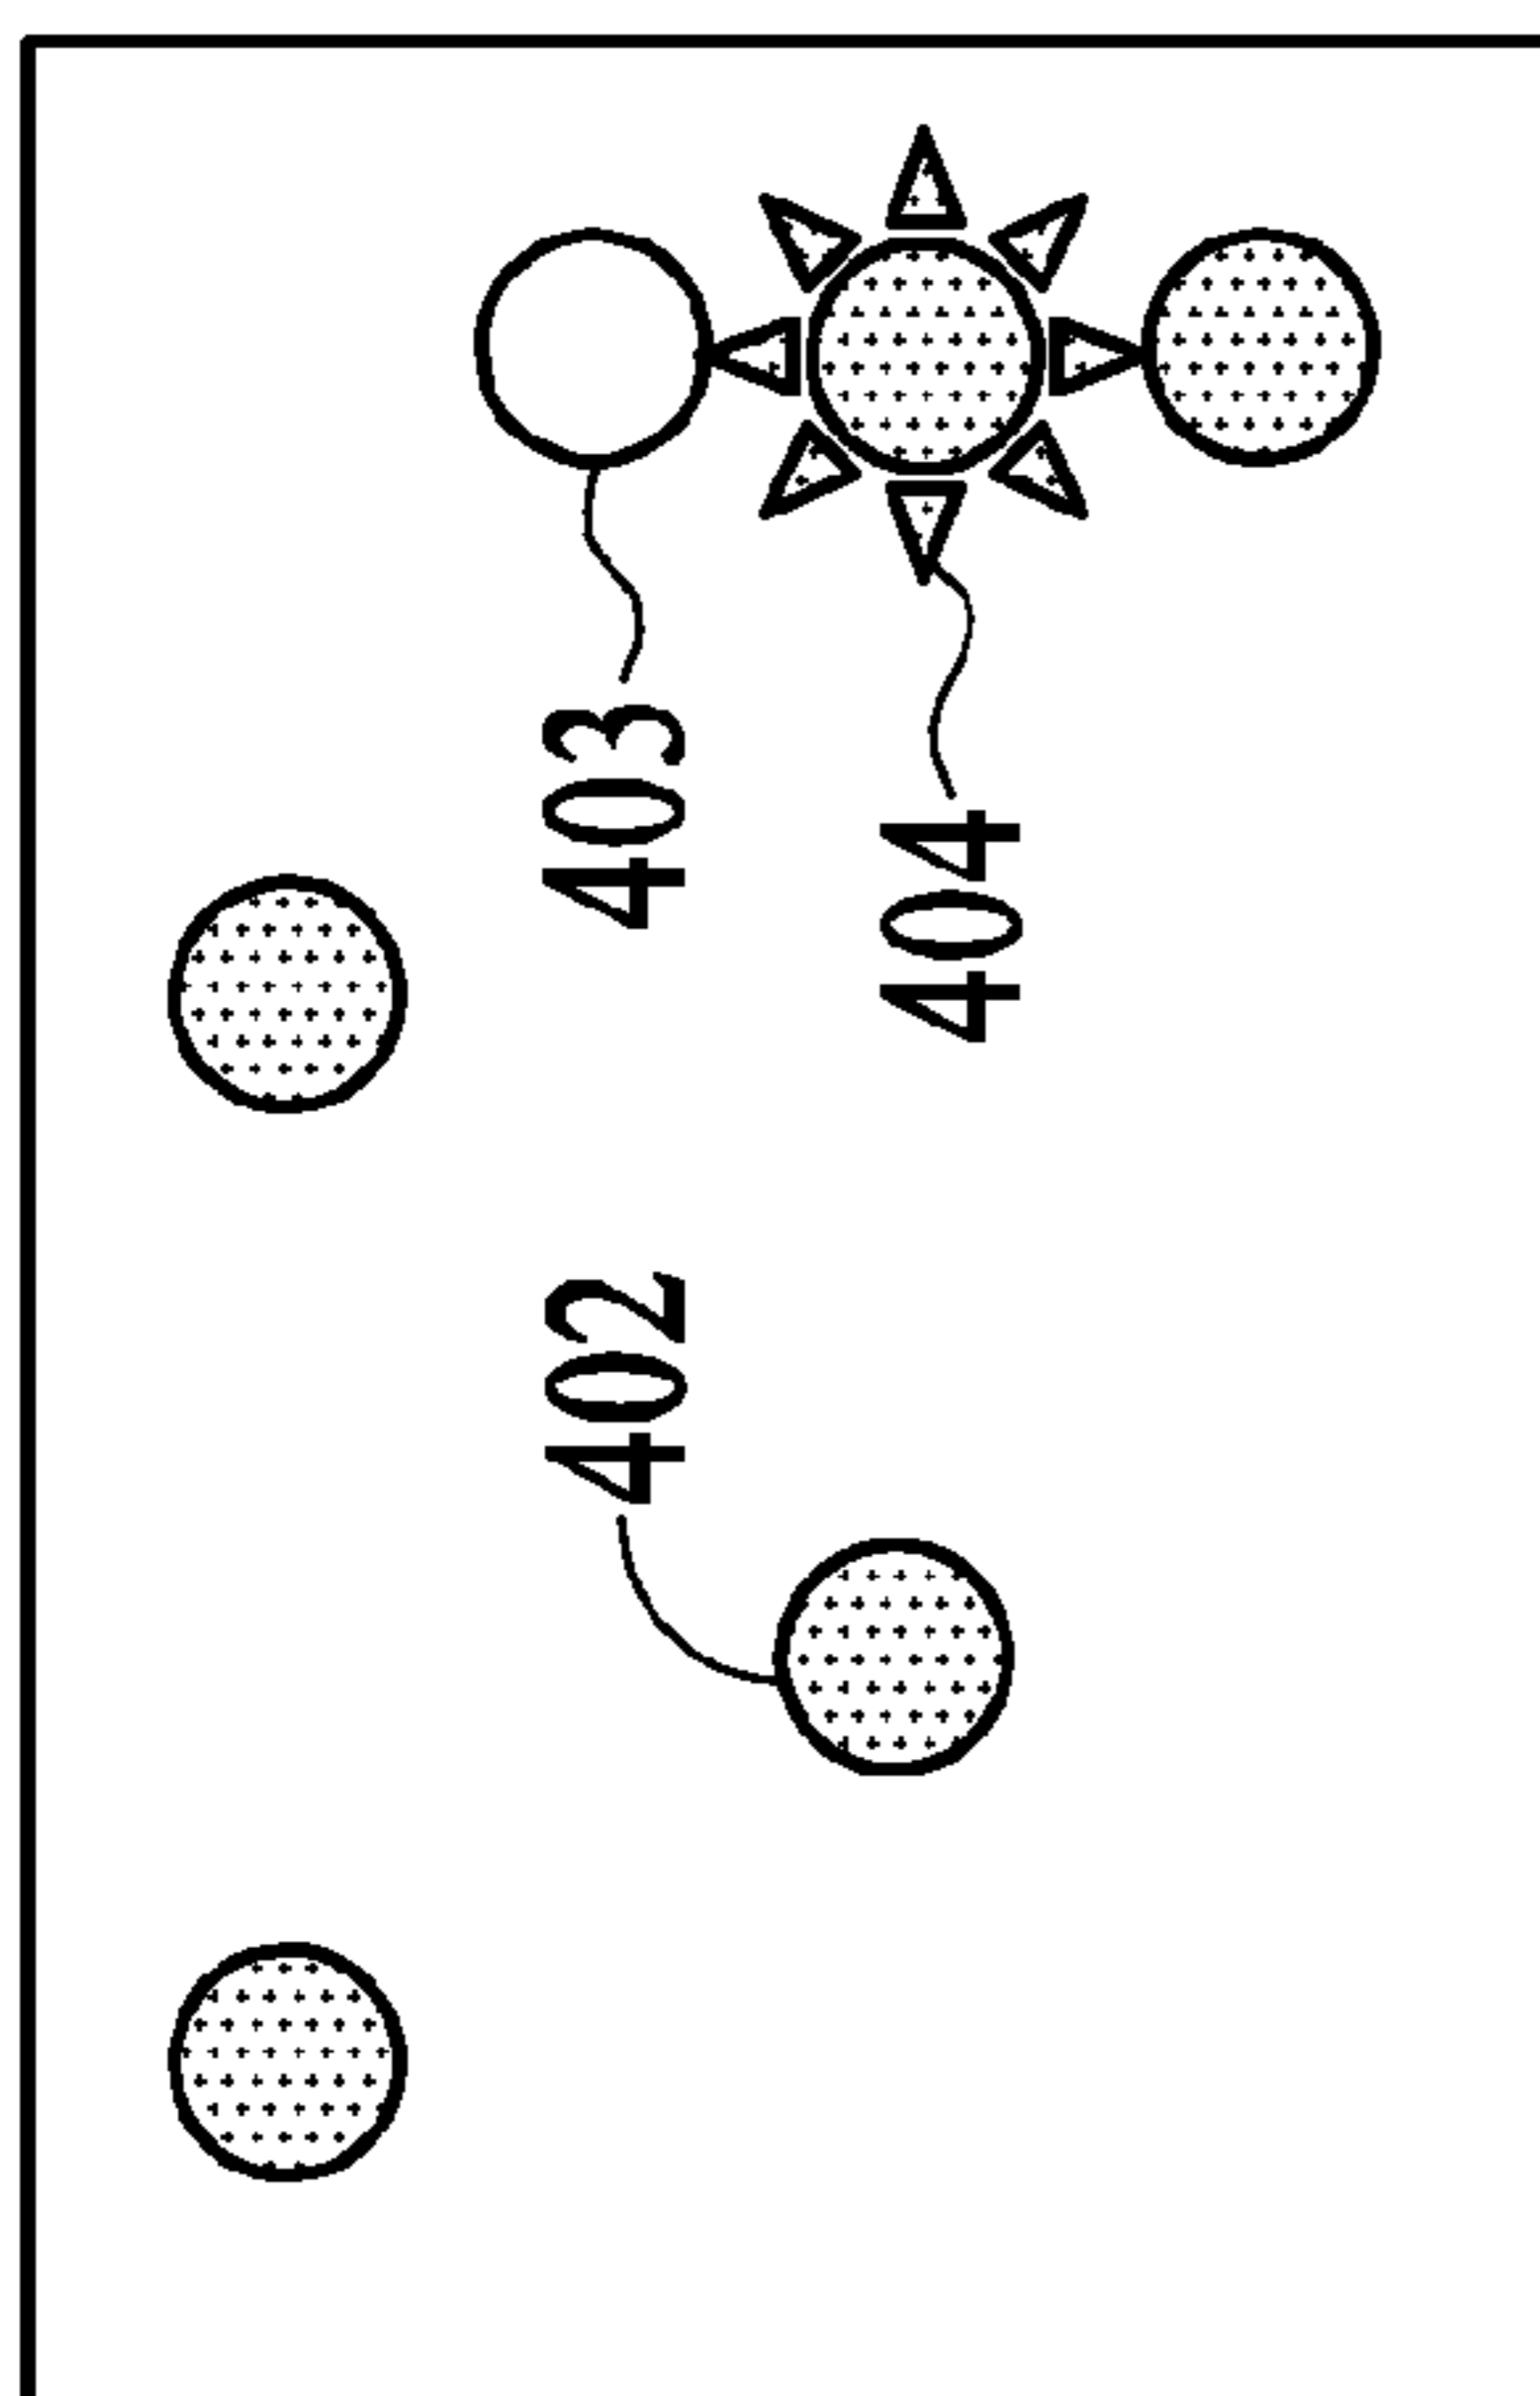


FIG. 4A

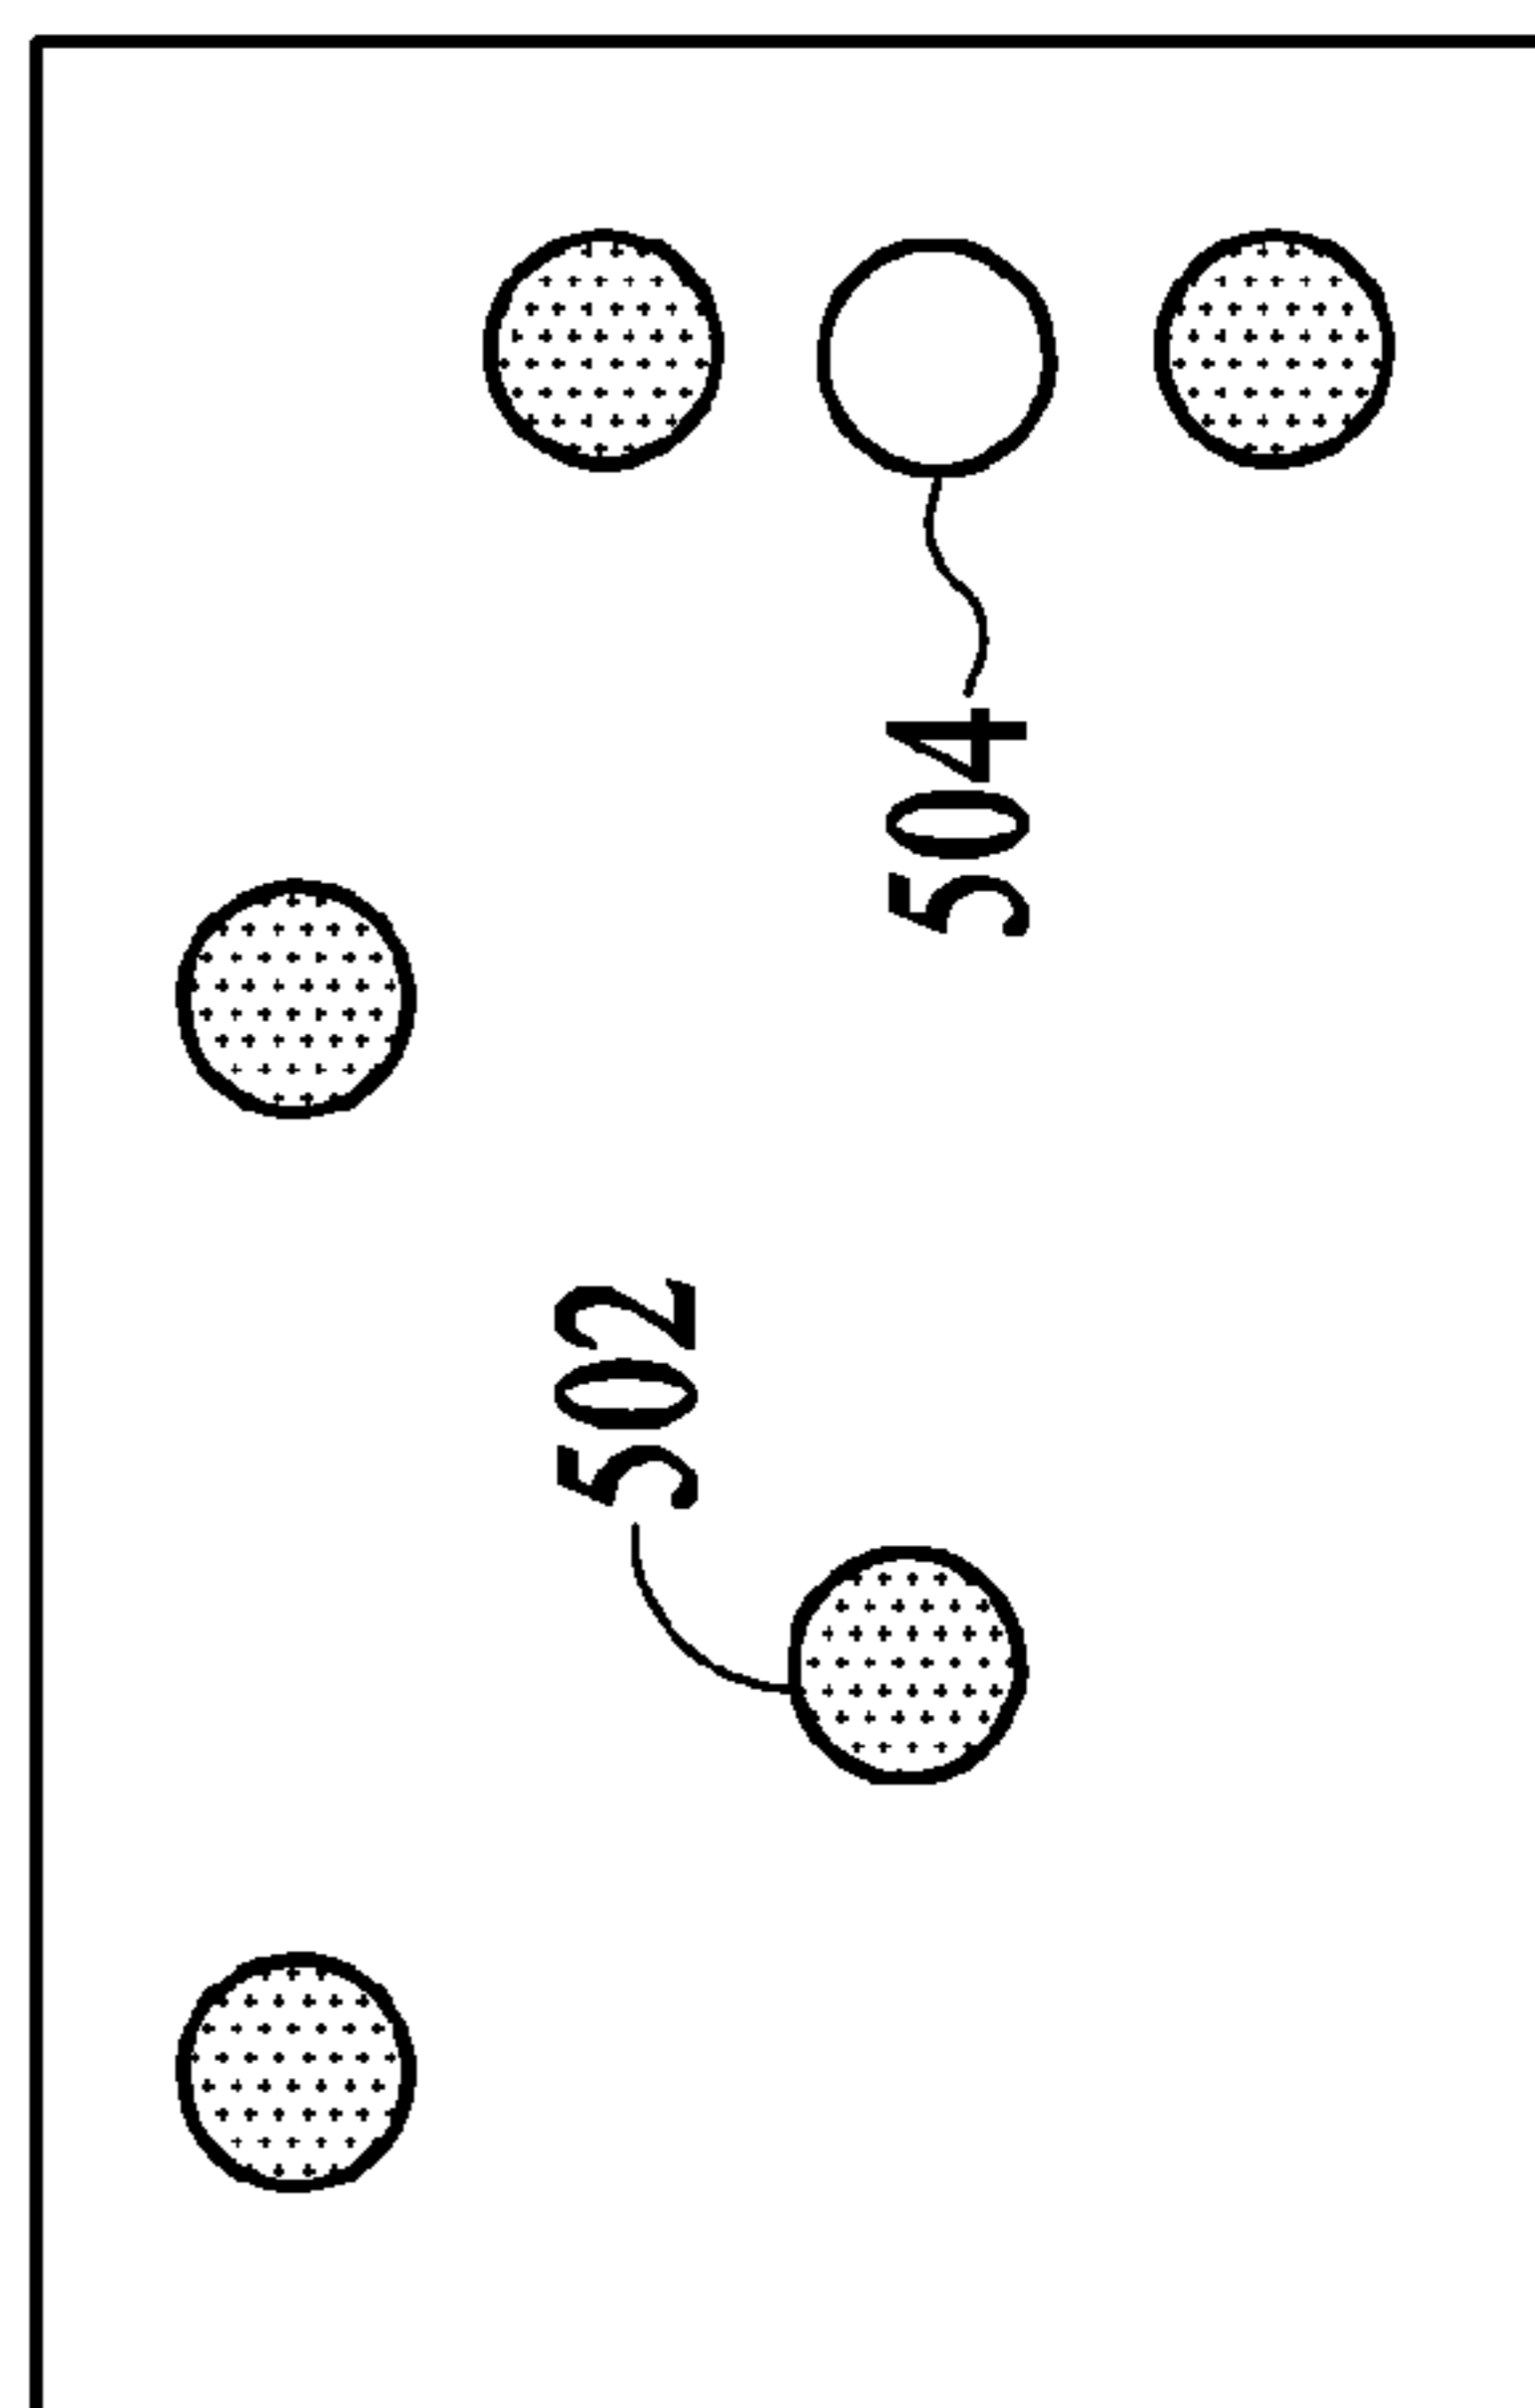


FIG. 5A

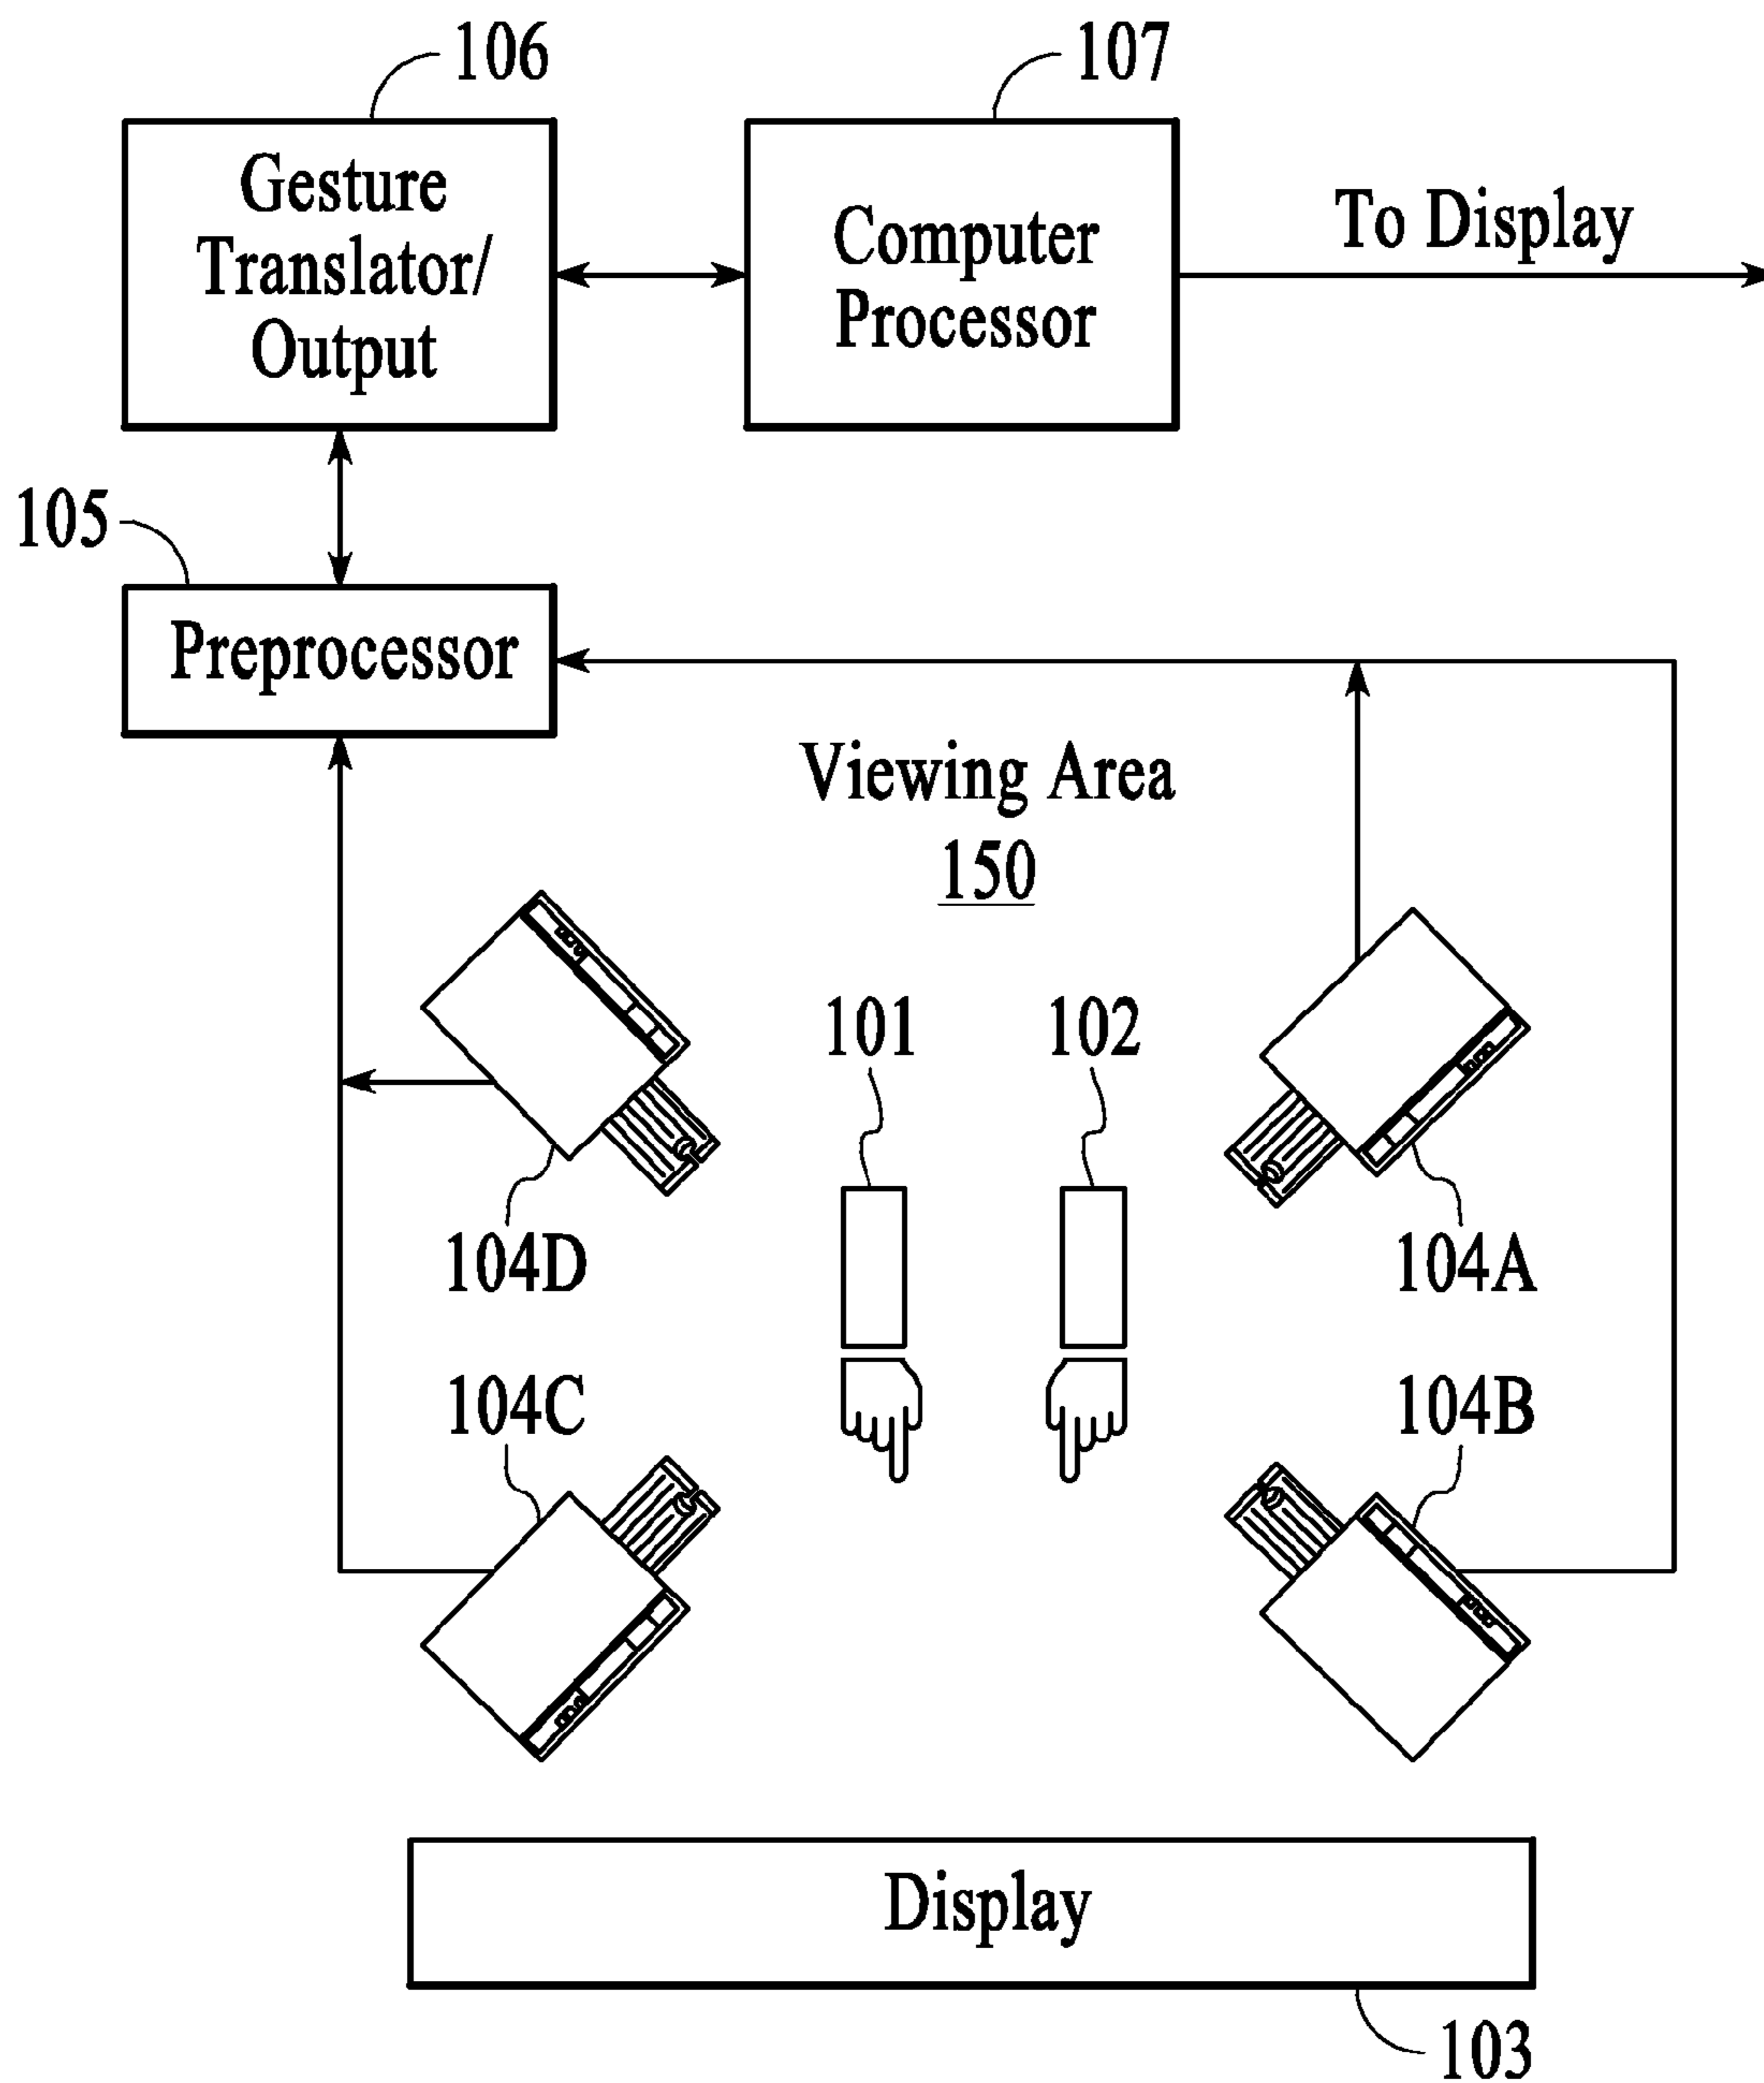


FIG. 6

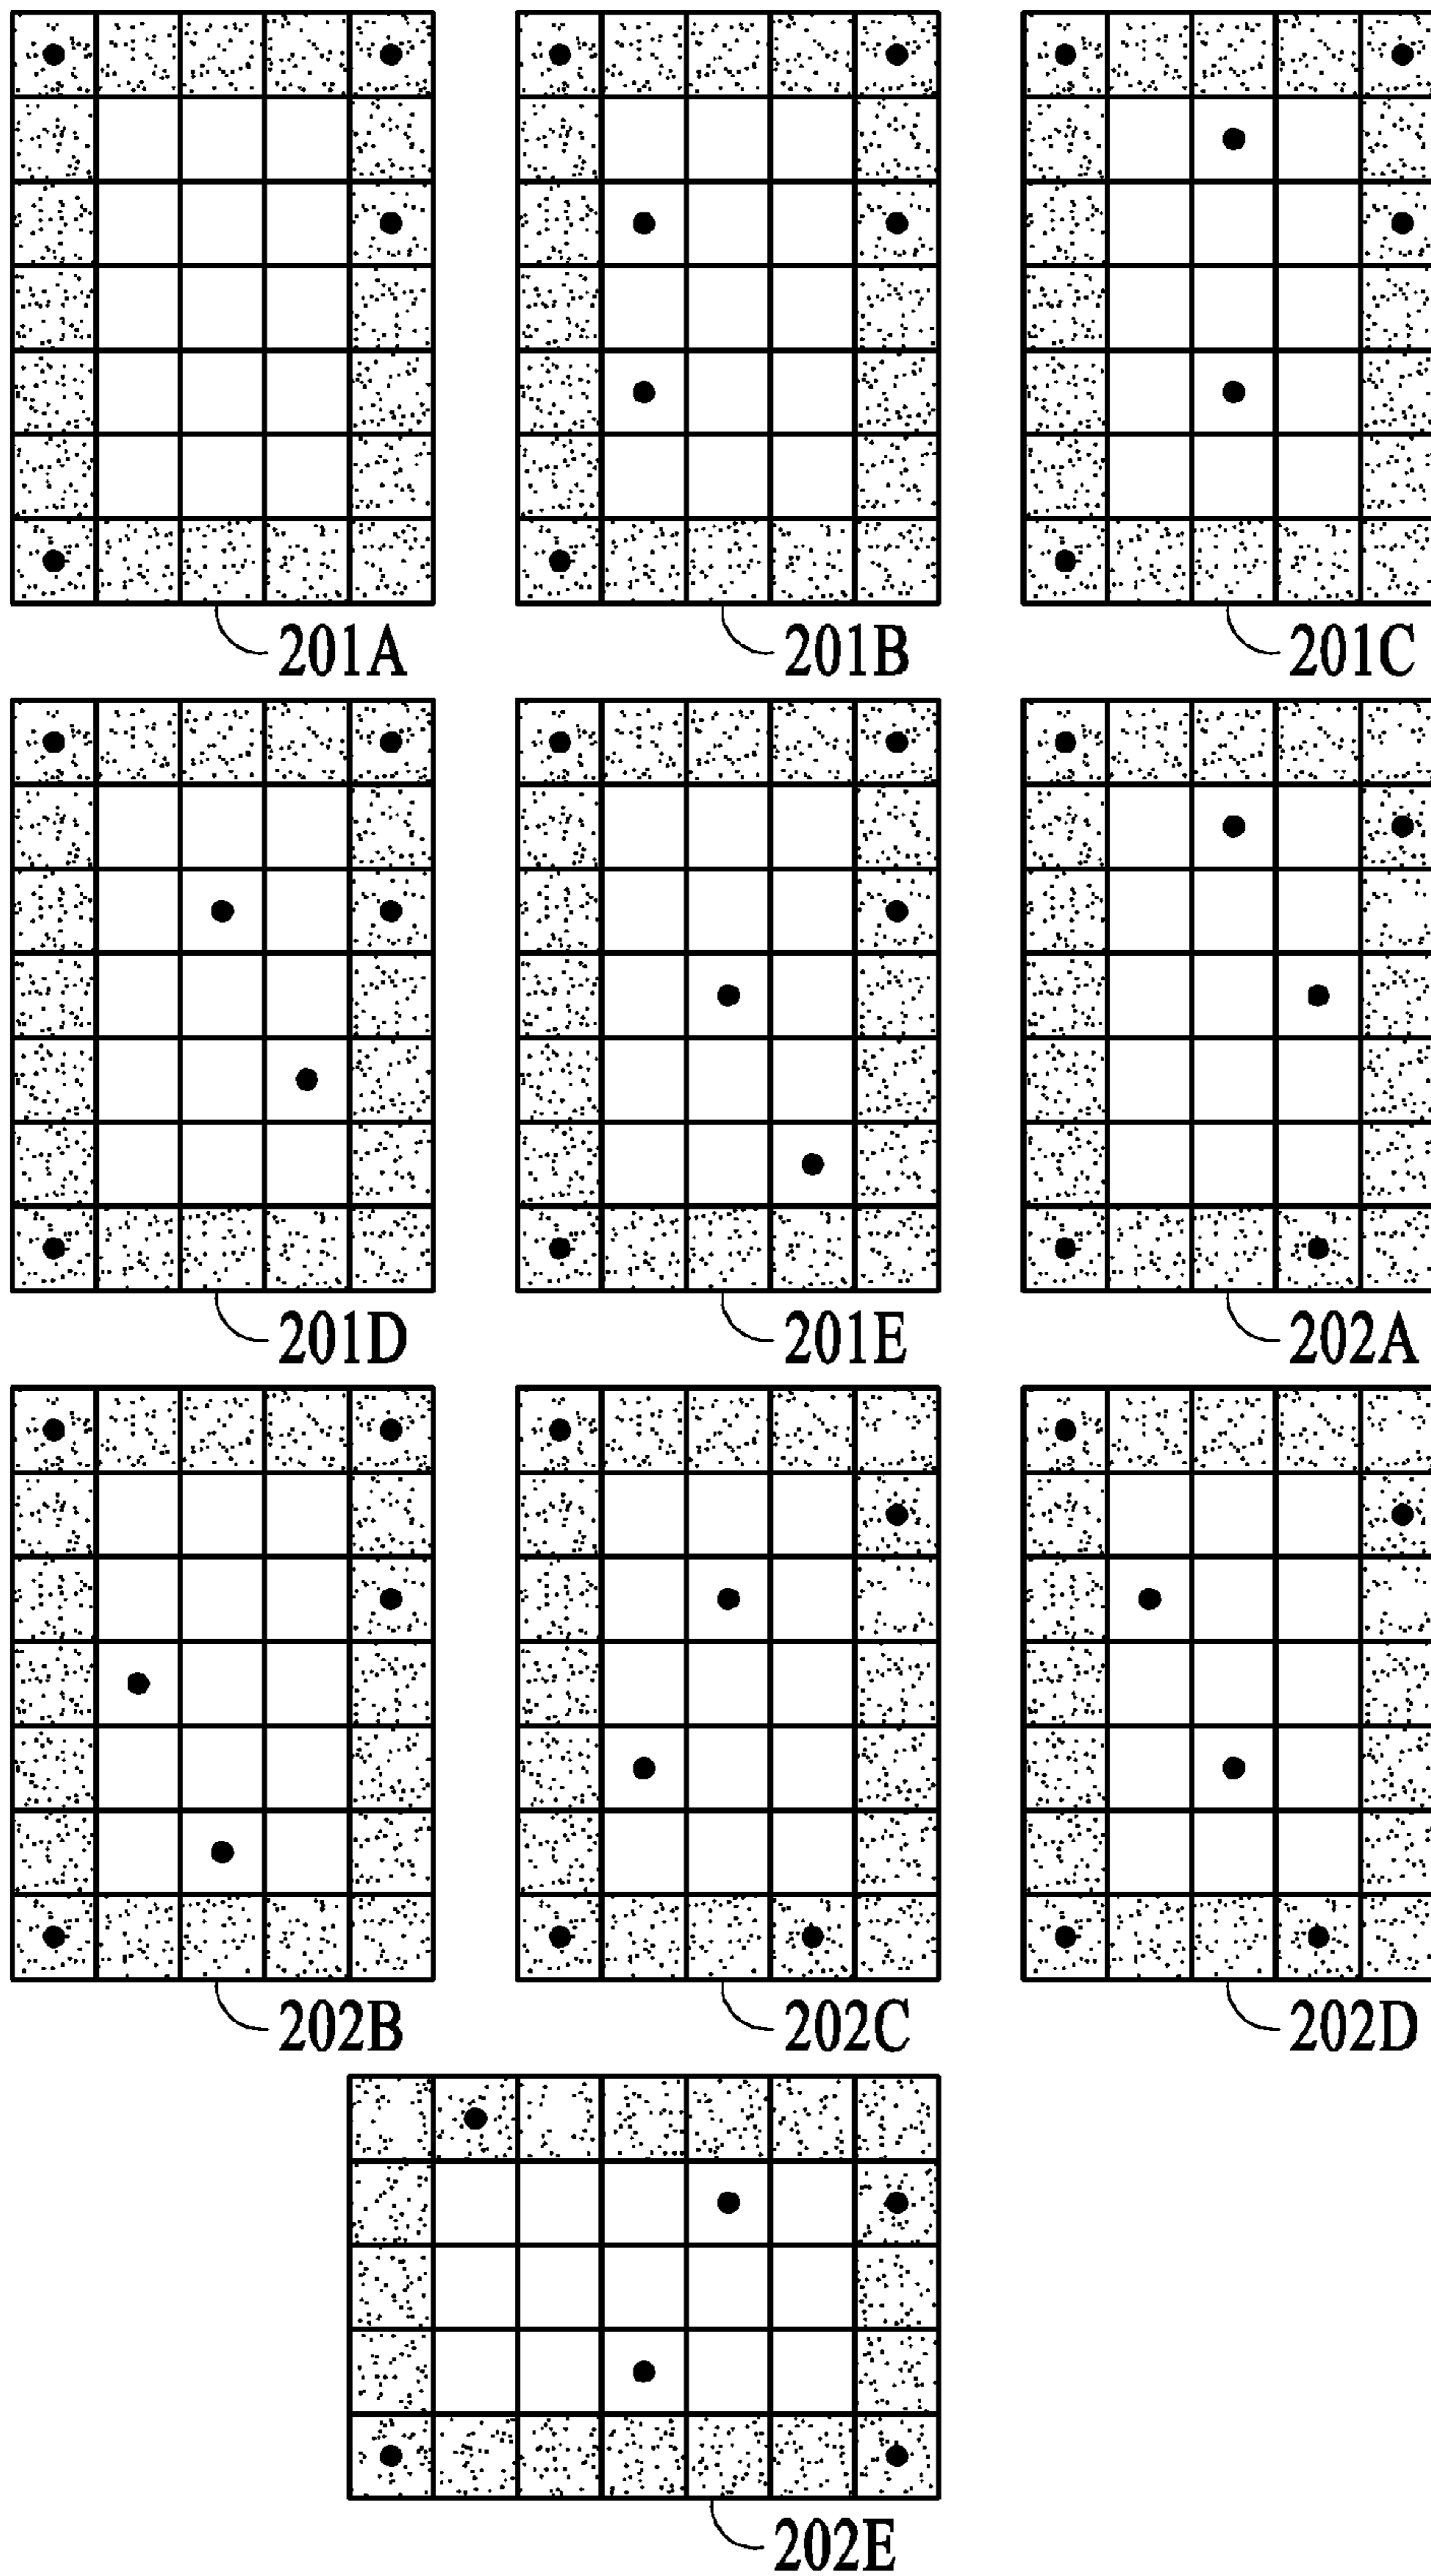


FIG. 7

1. Depict pose with left hand as viewed from back

p = pinkie finger
 r = ring finger
 m = middle finger
 i = index finger
 t = thumb

^ = curled non-thumb
 > = curled thumb
 | = straight finger or thumb pointed straight up
 \ or / = straight finger or thumb pointed at angle
 - = thumb pointing straight sideways
 x = finger or thumb pointing into plane

Pose name	p	r	m	i	t
flat					
fist	^	^	^	^	>
mime gun	^	^	^		-
2 or peace	^	^	\	/	>
one-finger point	^	^	^		>
two-finger point	^	^			>
x-y-z	^	^	x		-
ok				^	>
pinkie point		^	^	^	>
bracket	x	x	x	x	x
4	\	\		/	>
3	^	\		/	>
5	\	\		/	/

FIG.8

2. Add hand orientation to complete pose

must specify two variables:

1. palm direction (if hand were flat)
2. finger direction (if hand were flat)

-	medial
+	lateral
x	anterior
*	posterior
^	cranial
v	caudal

orientation variables come after colon, e.g.:

^ ^ x - : - x	=	x·y·z start position
^ ^ \ / > : * v	=	upside-down v

FIG.9

3. Two-hand combos

Hand 1	Hand 2	Pose
^ ^ ^ ^ > : x ^	^ ^ ^ ^ > : x ^	full stop
^ ^ ^ - : x -	^ ^ ^ - : x ^	snapshot
: v x	: - x	rudder and throttle start position

FIG.10

4. Orientation blends
 Achieve variable blending by enclosing pairs
 e.g.:

: (vx) (x^)	flat at 45 degrees pitch toward screen
^ ^ > : (- (- v)) x	two-finger point rolled medially to 22.5 degrees (halfway between palm medial and palm rolled to 45 degrees)

FIG.11

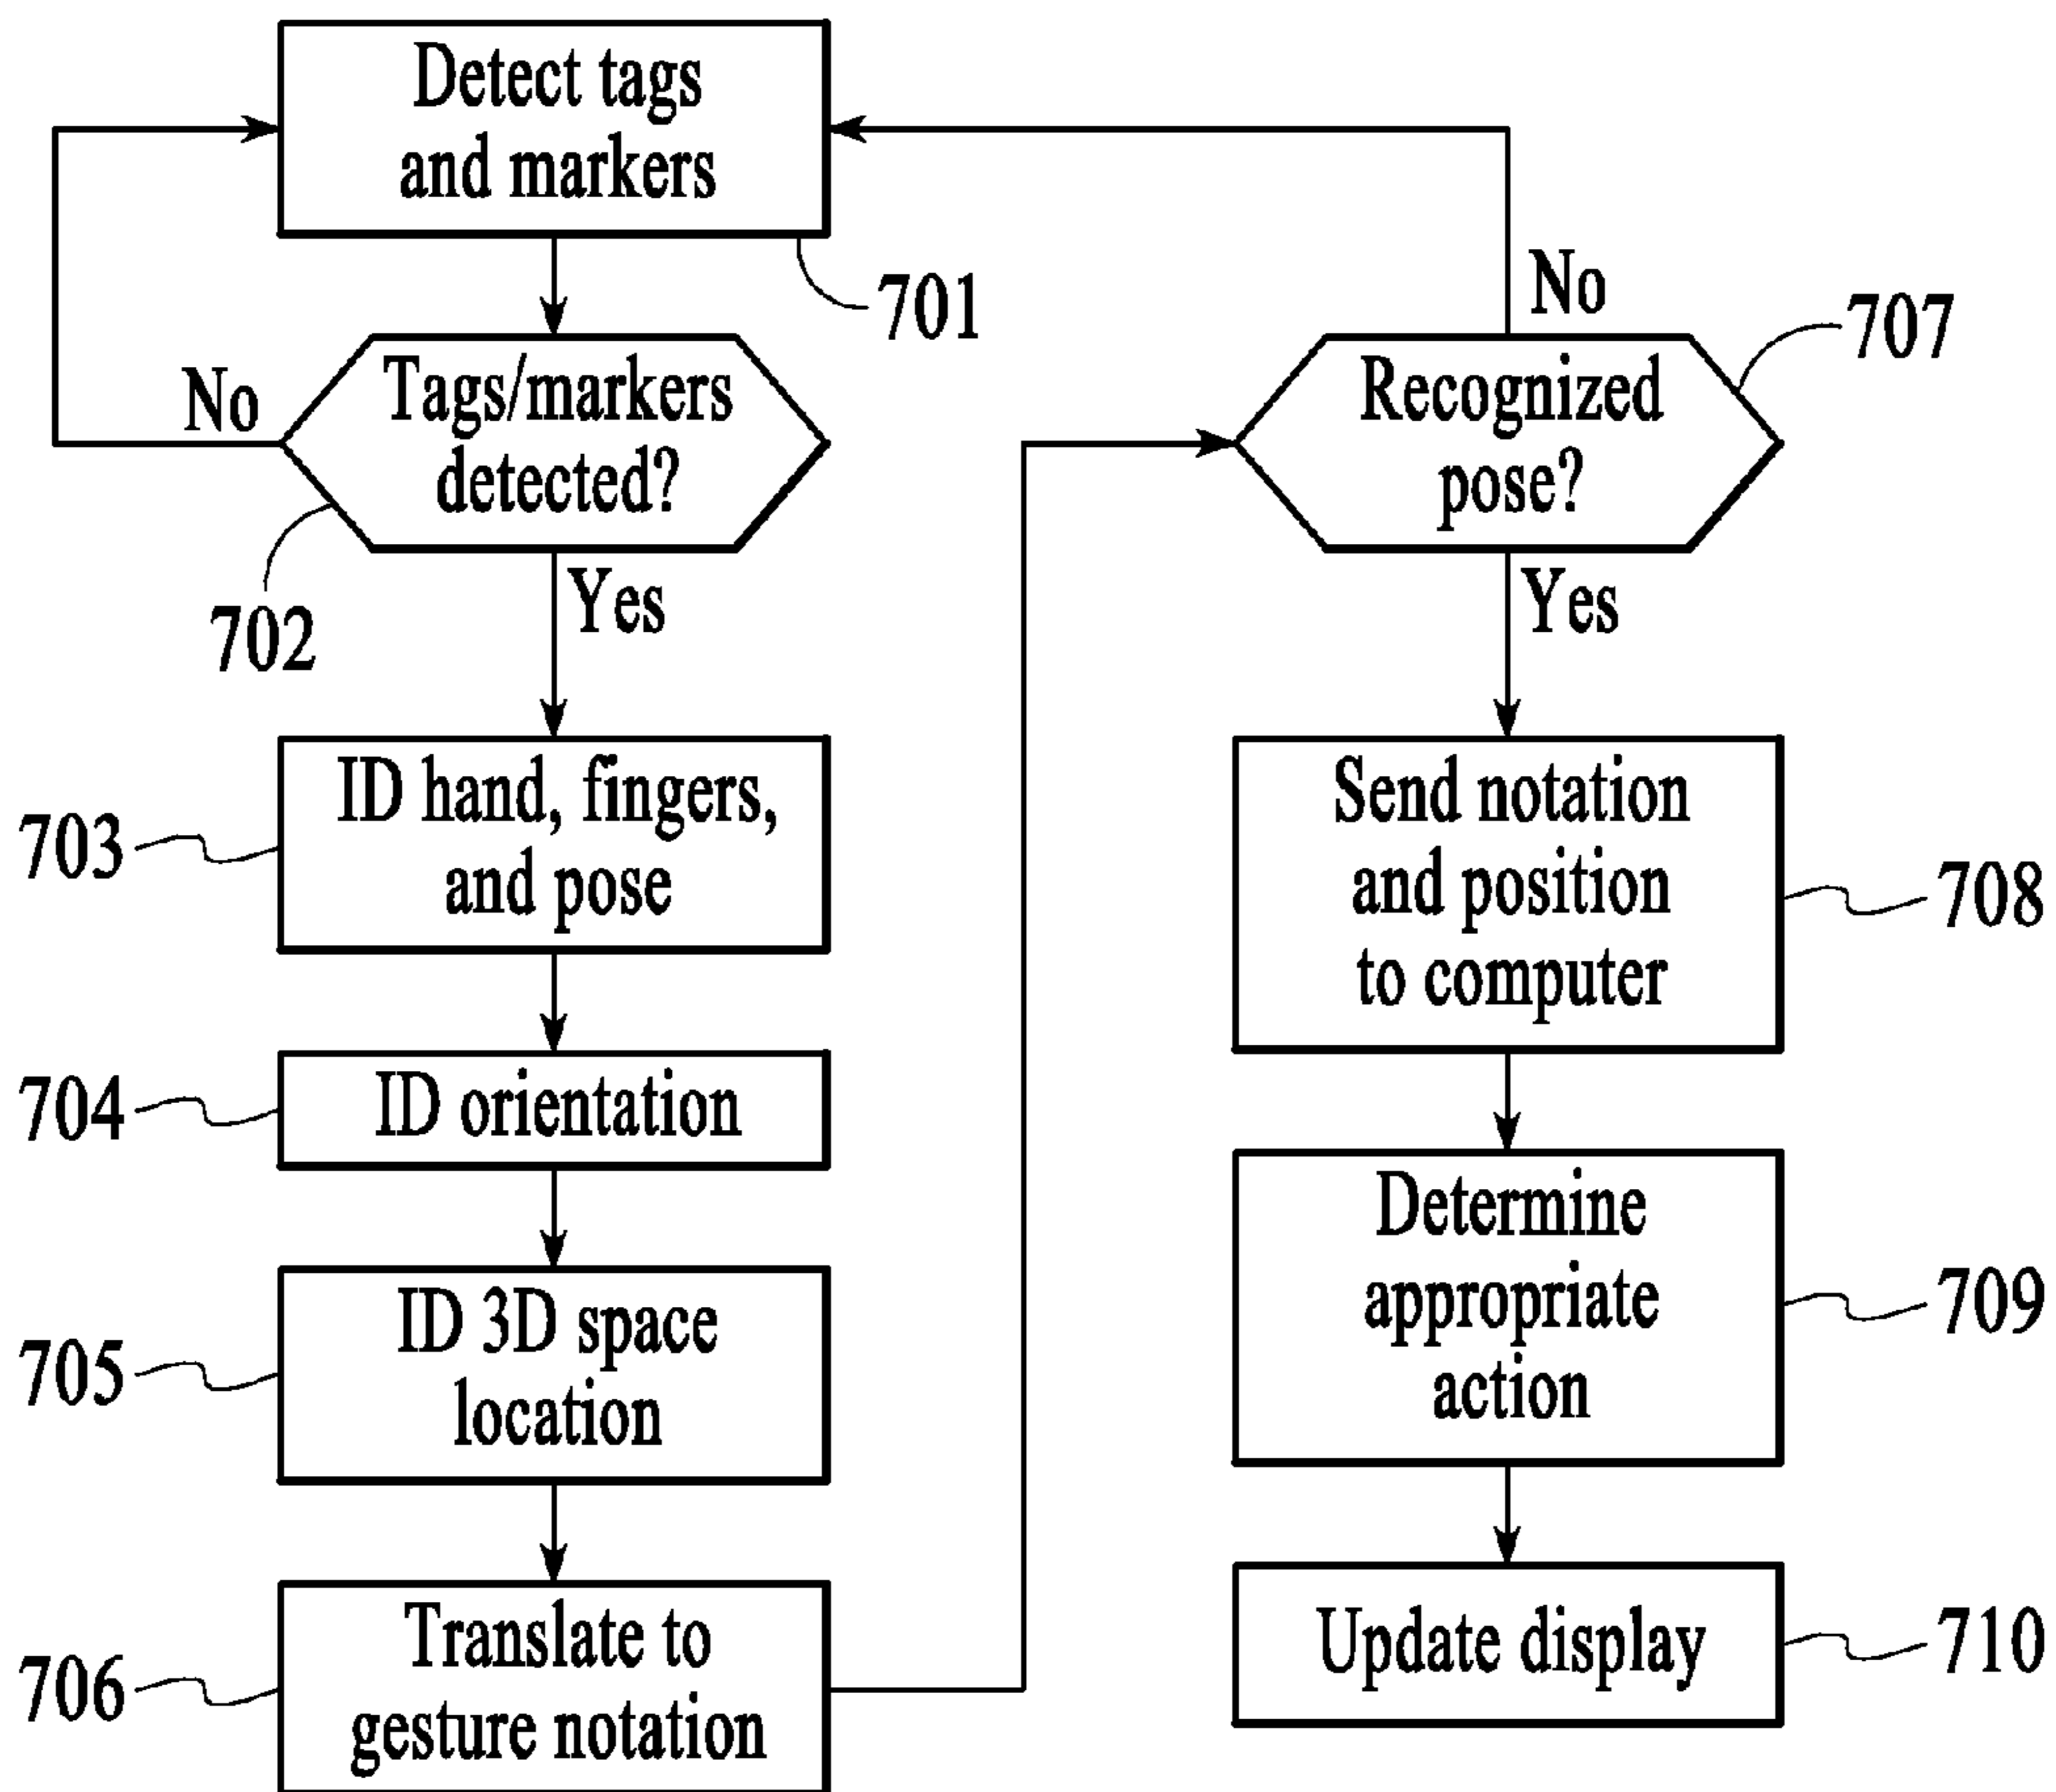


FIG.12

Gest I.D.	Description	Hand 1		Hand 2	
		Pose	Motion	Pose	Motion
1	point at object (invoke and move cursor)	^ ^ ^ - : - X	point mime gun		
2	select object	^ ^ ^ : - X	drop thumb to select		
3	move spatially / zoom in/out	^ ^ X - : - X	rotate/translate		
4	snapshot	^ ^ ^ - : X -	make square with 2 hands	^ ^ ^ - : X ^	make square with 2 hands
5	demarkate rectangular region	^ ^ ^ - : X -	make square then adjust size	^ ^ ^ - : X ^	make square then adjust size
6	clear the decks	: + X	sweep hand laterally	: - X	sweep hand medially
7	organize objects into a circle	^ ^ ^ - : - ^	look through circle of O.K. sign		
8	two-finger point at objects	^ ^ - : - X	point		
9	two-finger select object	^ ^ ^ : - X	drop thumb to select		
10	mark start time	X X X X X : - ^	strike pose		
11	mode change 1	: - ^	strike pose - make "T" with two hands	: v -	strike pose - make "T" with two hands
12	mode change 11	: - ^	strike pose - parallel hands	: - ^	strike pose - parallel hands
13	push back and slide workspace	- : X ^	push palm toward screen - - move sideways to find new regions		

FIG.13/1

14	enter sub-application		: x ^	strike pose		: x ^	strike pose
15	return from sub-application		: . ^	strike pose		: . ^	strike pose
16	select option	^ ^ ^	- : - x	medial roll			
17	roll time forward/back		: v x	Yaw hand at elbow while keeping hand parallel to floor			
18	stop time		: x ^	strike pose			
19	loop time	^ ^ ^	- : x ^	circular motion with "L"			
20	demarcate irregular region	^ ^ ^	- : v x	start with 2 finger tips together. 1 hand holds start position.	^ ^ ^	- : - x	other hand traces out shape - select "click" for vertices
21	tag object	^ ^ ^ >	- : x	pinky-point at object then roll hand medially			
22	group data streams	^ ^ ^	- : v x	bring finger tips of two hands together	^ ^ ^	- : v x	bring finger tips of two hands together
23	restore encapsulated workspace		: + x	sweep hand medially		: - x	sweep hand laterally

FIG.13/2

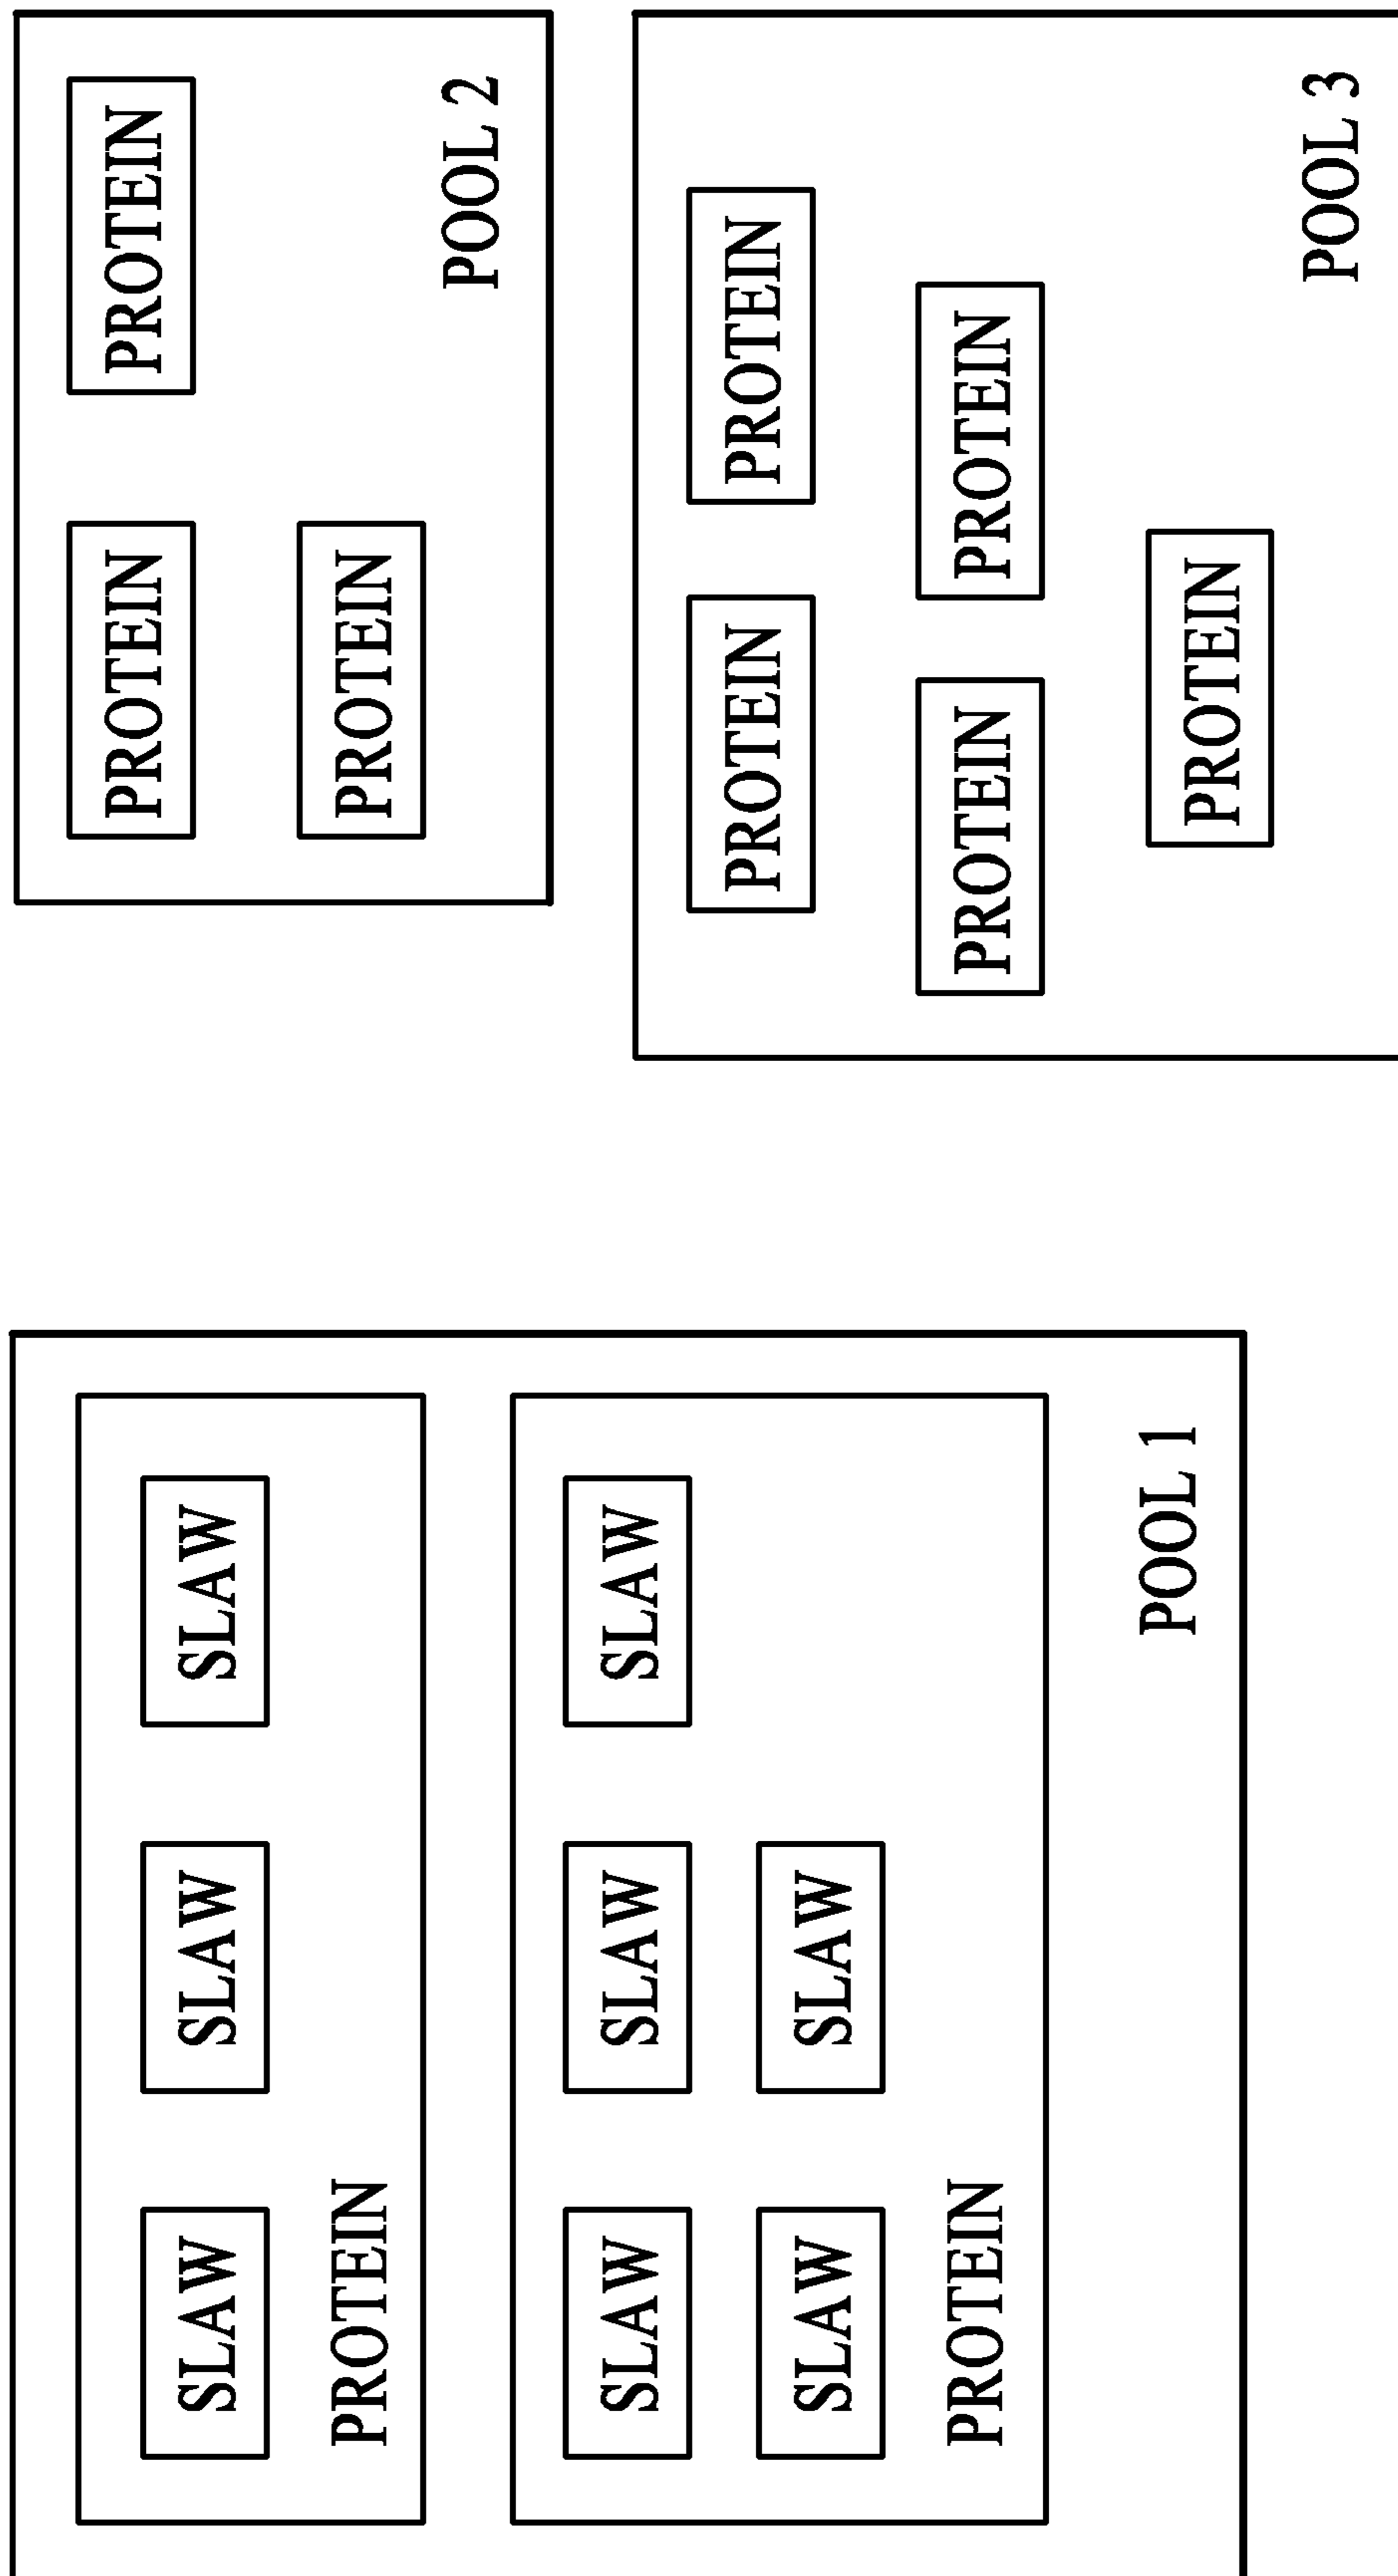


FIG. 14

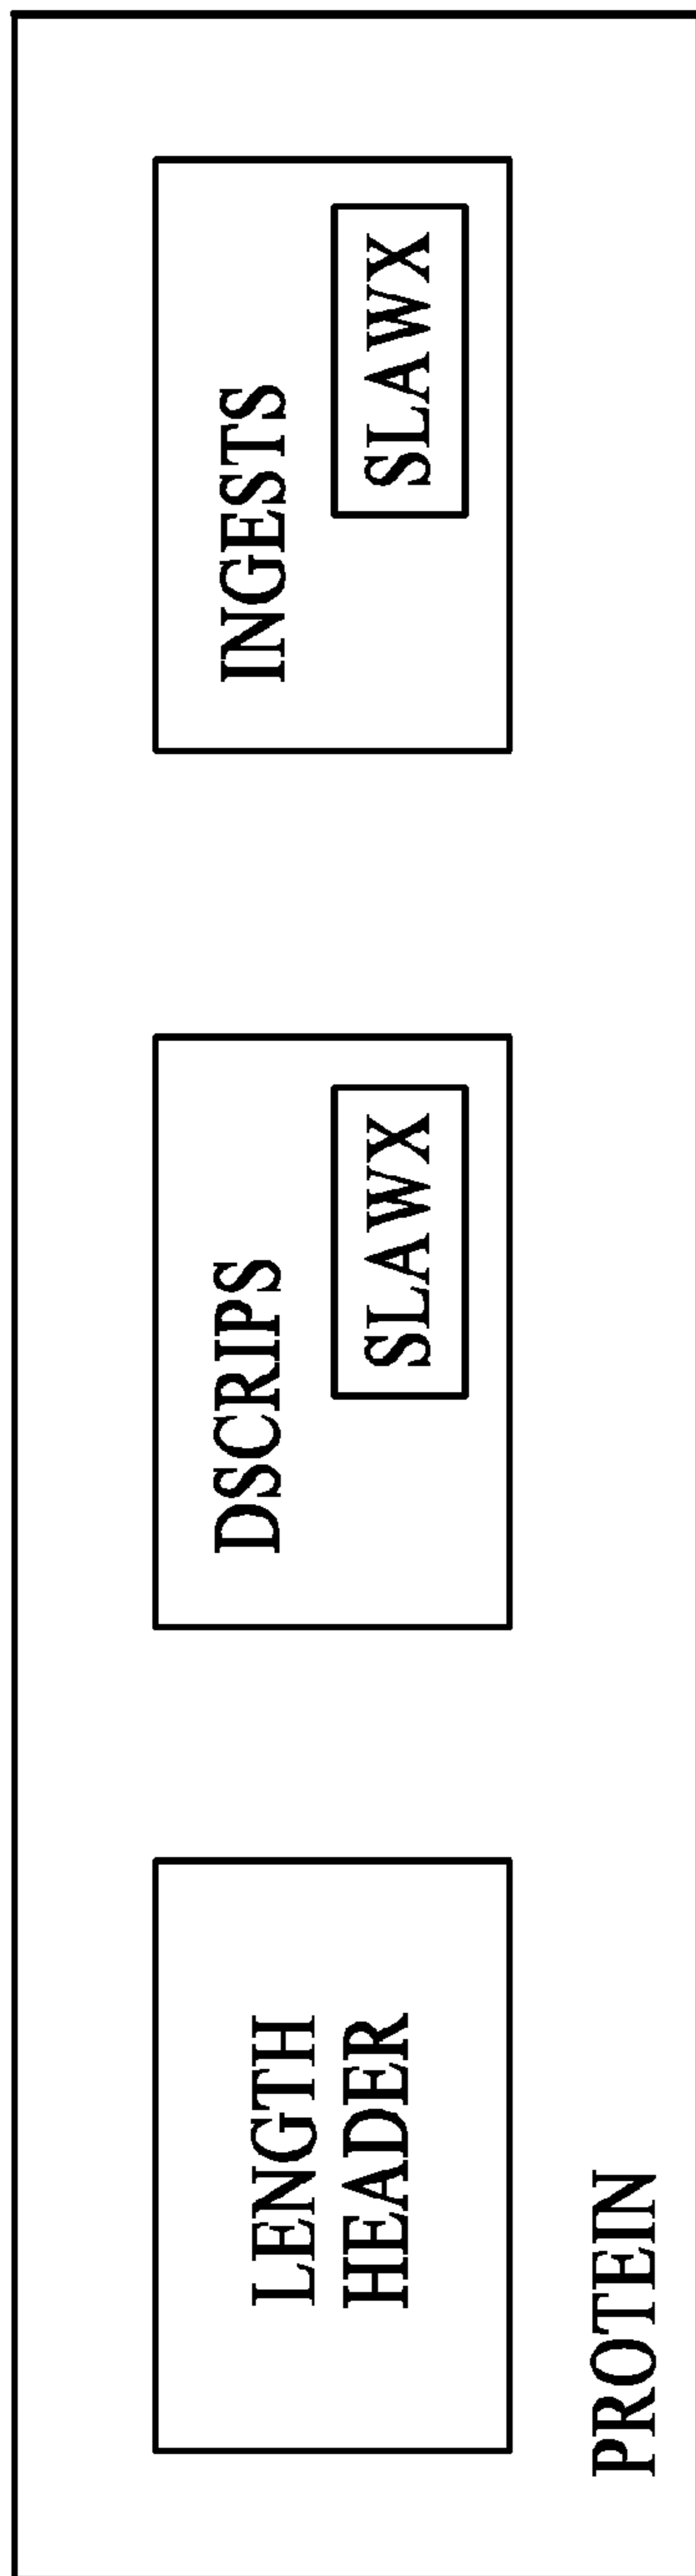


FIG.15

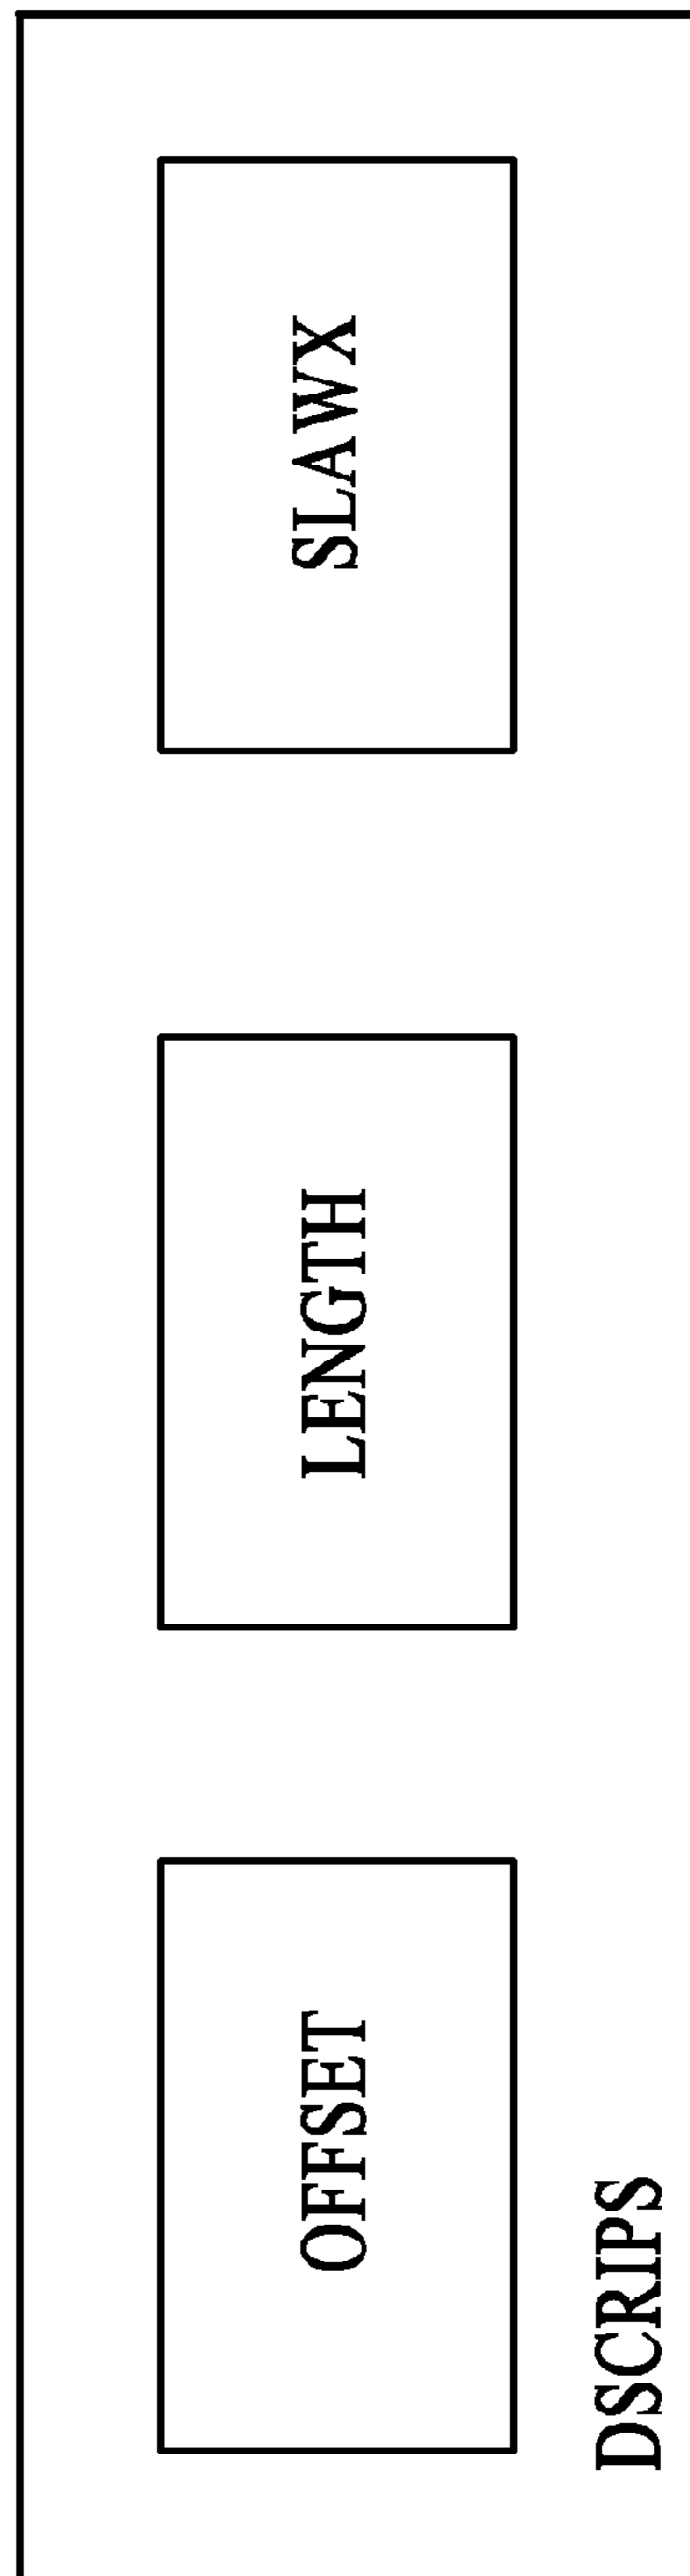


FIG.16

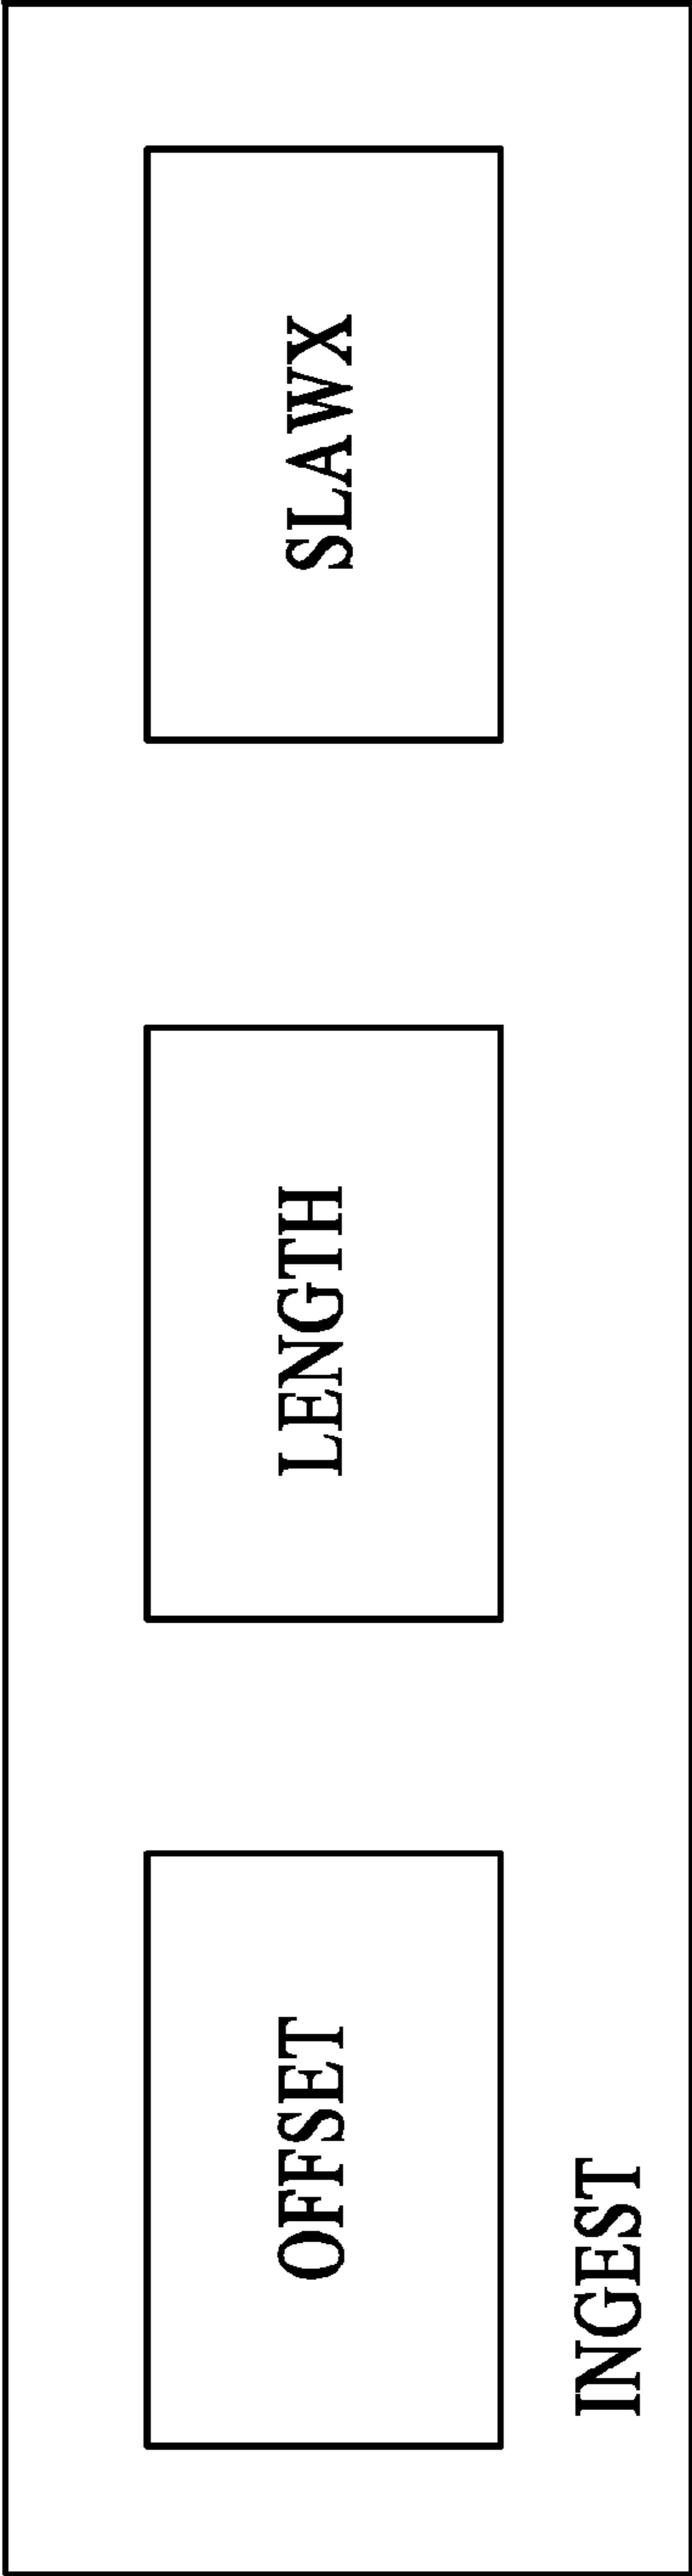


FIG.17

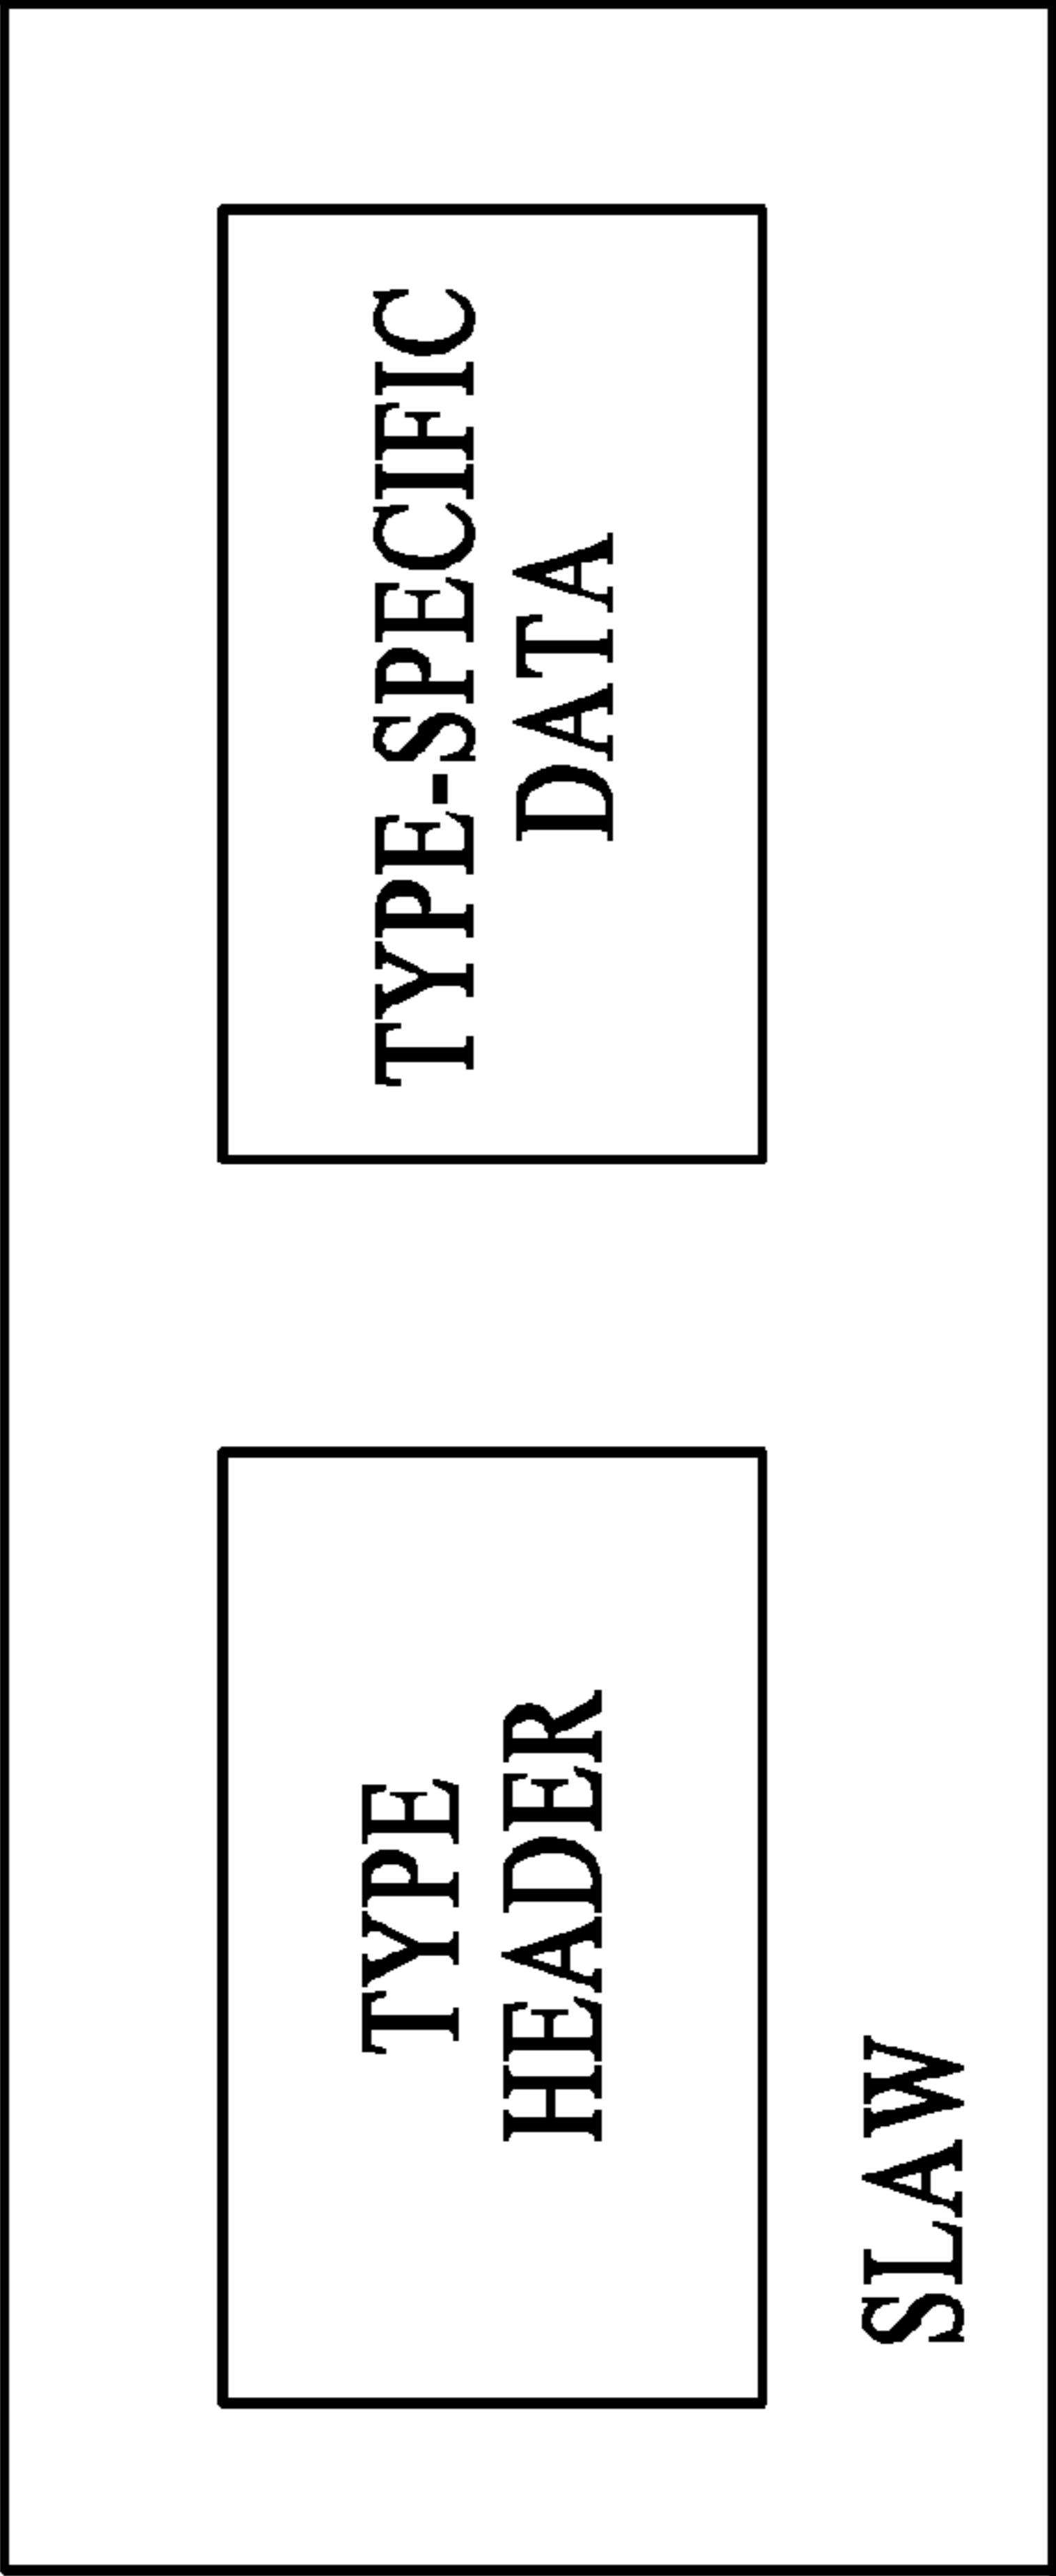


FIG.18

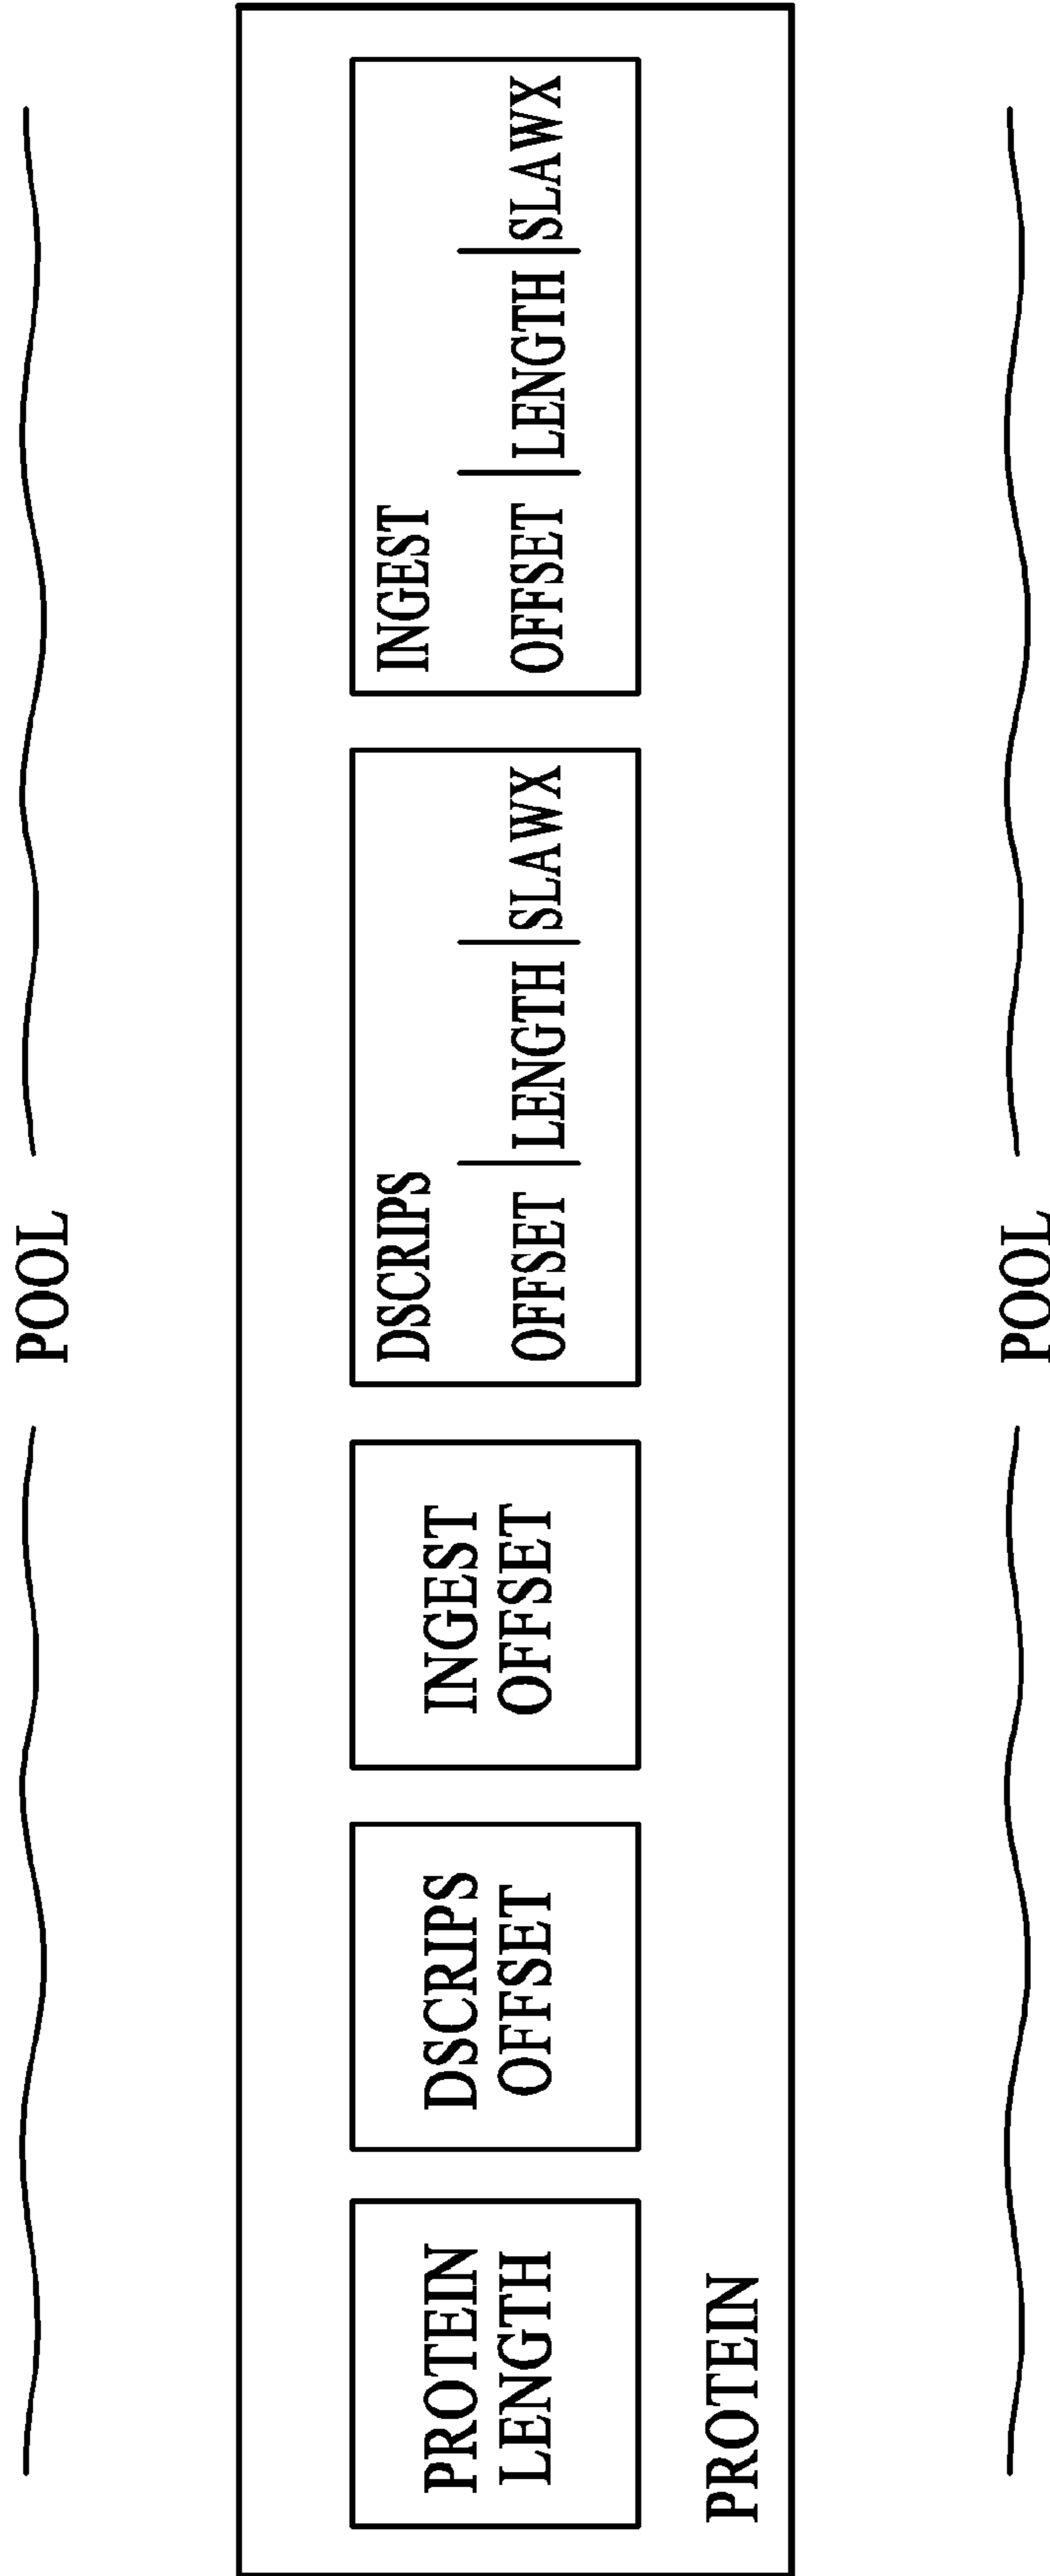


FIG.19A

first quadword of every slaw

	76543210	76543210	76543210	76543210
length-follows:	1xxxxxxx	xxxxxxx	xxxxxxx	xxxxxxx
eight-byte length:	11xxxxxx	xxxxxxx	xxxxxxx	xxxxxxx
wee cons:	01xxxxxx	xxxxxxx	xxxxxxx	xxxxxxx
wee cons quadlen:	rrqqqqqq	qqqqqqqq	qqqqqqqq	qqqqqqqq
wee string:	001xxxxx	xxxxxxx	xxxxxxx	xxxxxxx
wee string quadlen:	rrrqqqqq	qqqqqqqq	qqqqqqqq	qqqqqqqq
wee list:	0001xxxx	xxxxxxx	xxxxxxx	xxxxxxx
wee list quadlen:	rrrrqqqq	qqqqqqqq	qqqqqqqq	qqqqqqqq
full string:	1*100000	00000000	00000000	00000001
full cons:	1*100000	00000000	00000000	00000010
full list:	1*100000	00000000	00000000	00000011
(the penulti-MSB above is zero or one as the length is contained in the next one or two quadwords, i.e. if it's a four or eight byte length, per the 'eight-byte length' bit description second from top)				
numeric:	00001xxx	xxxxxxx	xxxxxxx	xxxxxxx
numeric float:	xxxxx1xx	xxxxxxx	xxxxxxx	xxxxxxx
numeric complex:	xxxxxx1x	xxxxxxx	xxxxxxx	xxxxxxx
numeric unsigned:	xxxxxxx1	xxxxxxx	xxxxxxx	xxxxxxx
numeric wide:	xxxxxxx	1xxxxxx	xxxxxxx	xxxxxxx
numeric stumpy:	xxxxxxx	x1xxxxx	xxxxxxx	xxxxxxx
numeric reserved:	xxxxxxx	xx1xxxx	xxxxxxx	xxxxxxx

FIG.19B/1

(wide and stumpy conspire to express whether the number in question is 8, 16, 32, or 64 bits long; neither-wide-nor-stumpy, i.e. both zero, is sort of canonical and thus means 32 bits; stumpy alone is 8; stumpy and wide is 16; and just wide is 64)

numeric 2-vector:	xxxxxxxx	xxx01xxx	xxxxxxxx	xxxxxxxx
numeric 3-vector:	xxxxxxxx	xxx10xxx	xxxxxxxx	xxxxxxxx
numeric 4-vector:	xxxxxxxx	xxx11xxx	xxxxxxxx	xxxxxxxx

for any numeric entity, array or not, a size-in-bytes-minus-one is stored in the last eight bits -- if a singleton, this describes the size of the data part; if an array, it's the size of a single element -- so:

num'c unit bsize mask: 00001xxx xxxxxxxx xxxxxxxx mmmmmmmm

and for arrays, there're these:

num'c breadth follows:	xxxxxxxx	xxxxx1xx	xxxxxxxx	xxxxxxxx
num'c 8-byte breadth:	xxxxxxxx	xxxxx11x	xxxxxxxx	xxxxxxxx
num'c wee breadth mask:	xxxxxxxx	xxxxx0mm	mmmmmmmm	xxxxxxxx

FIG.19B/2

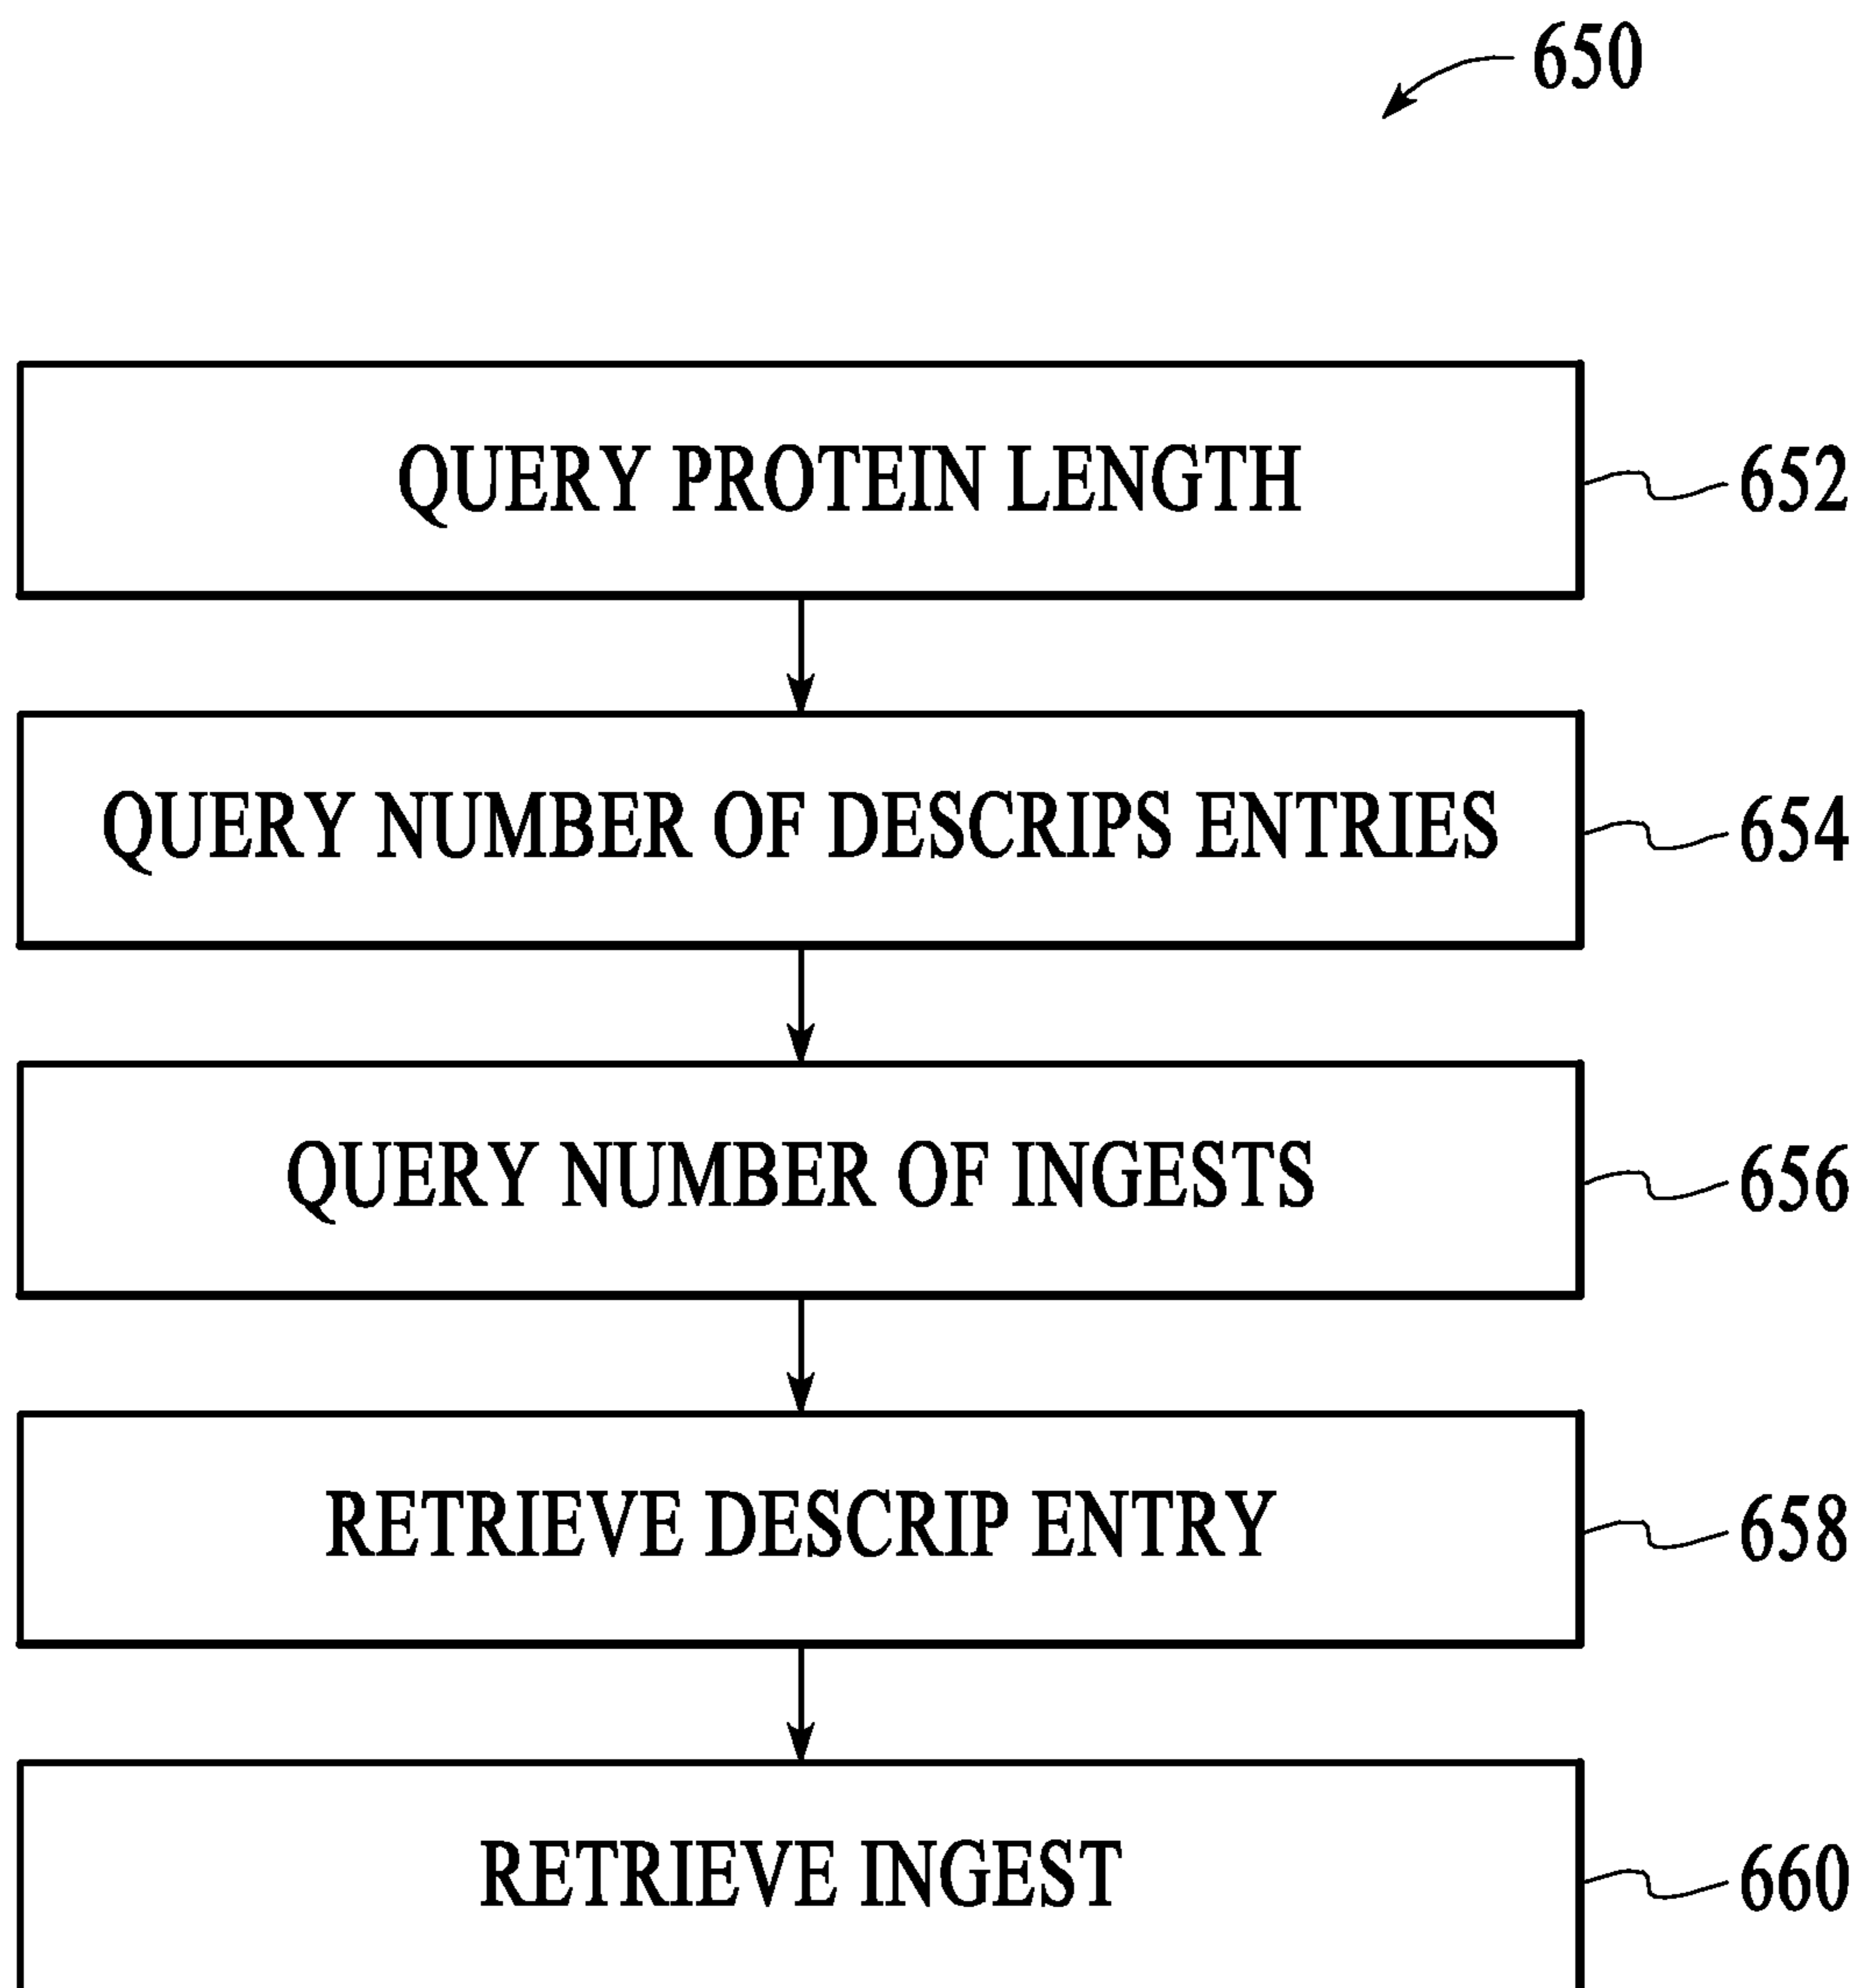


FIG.19C

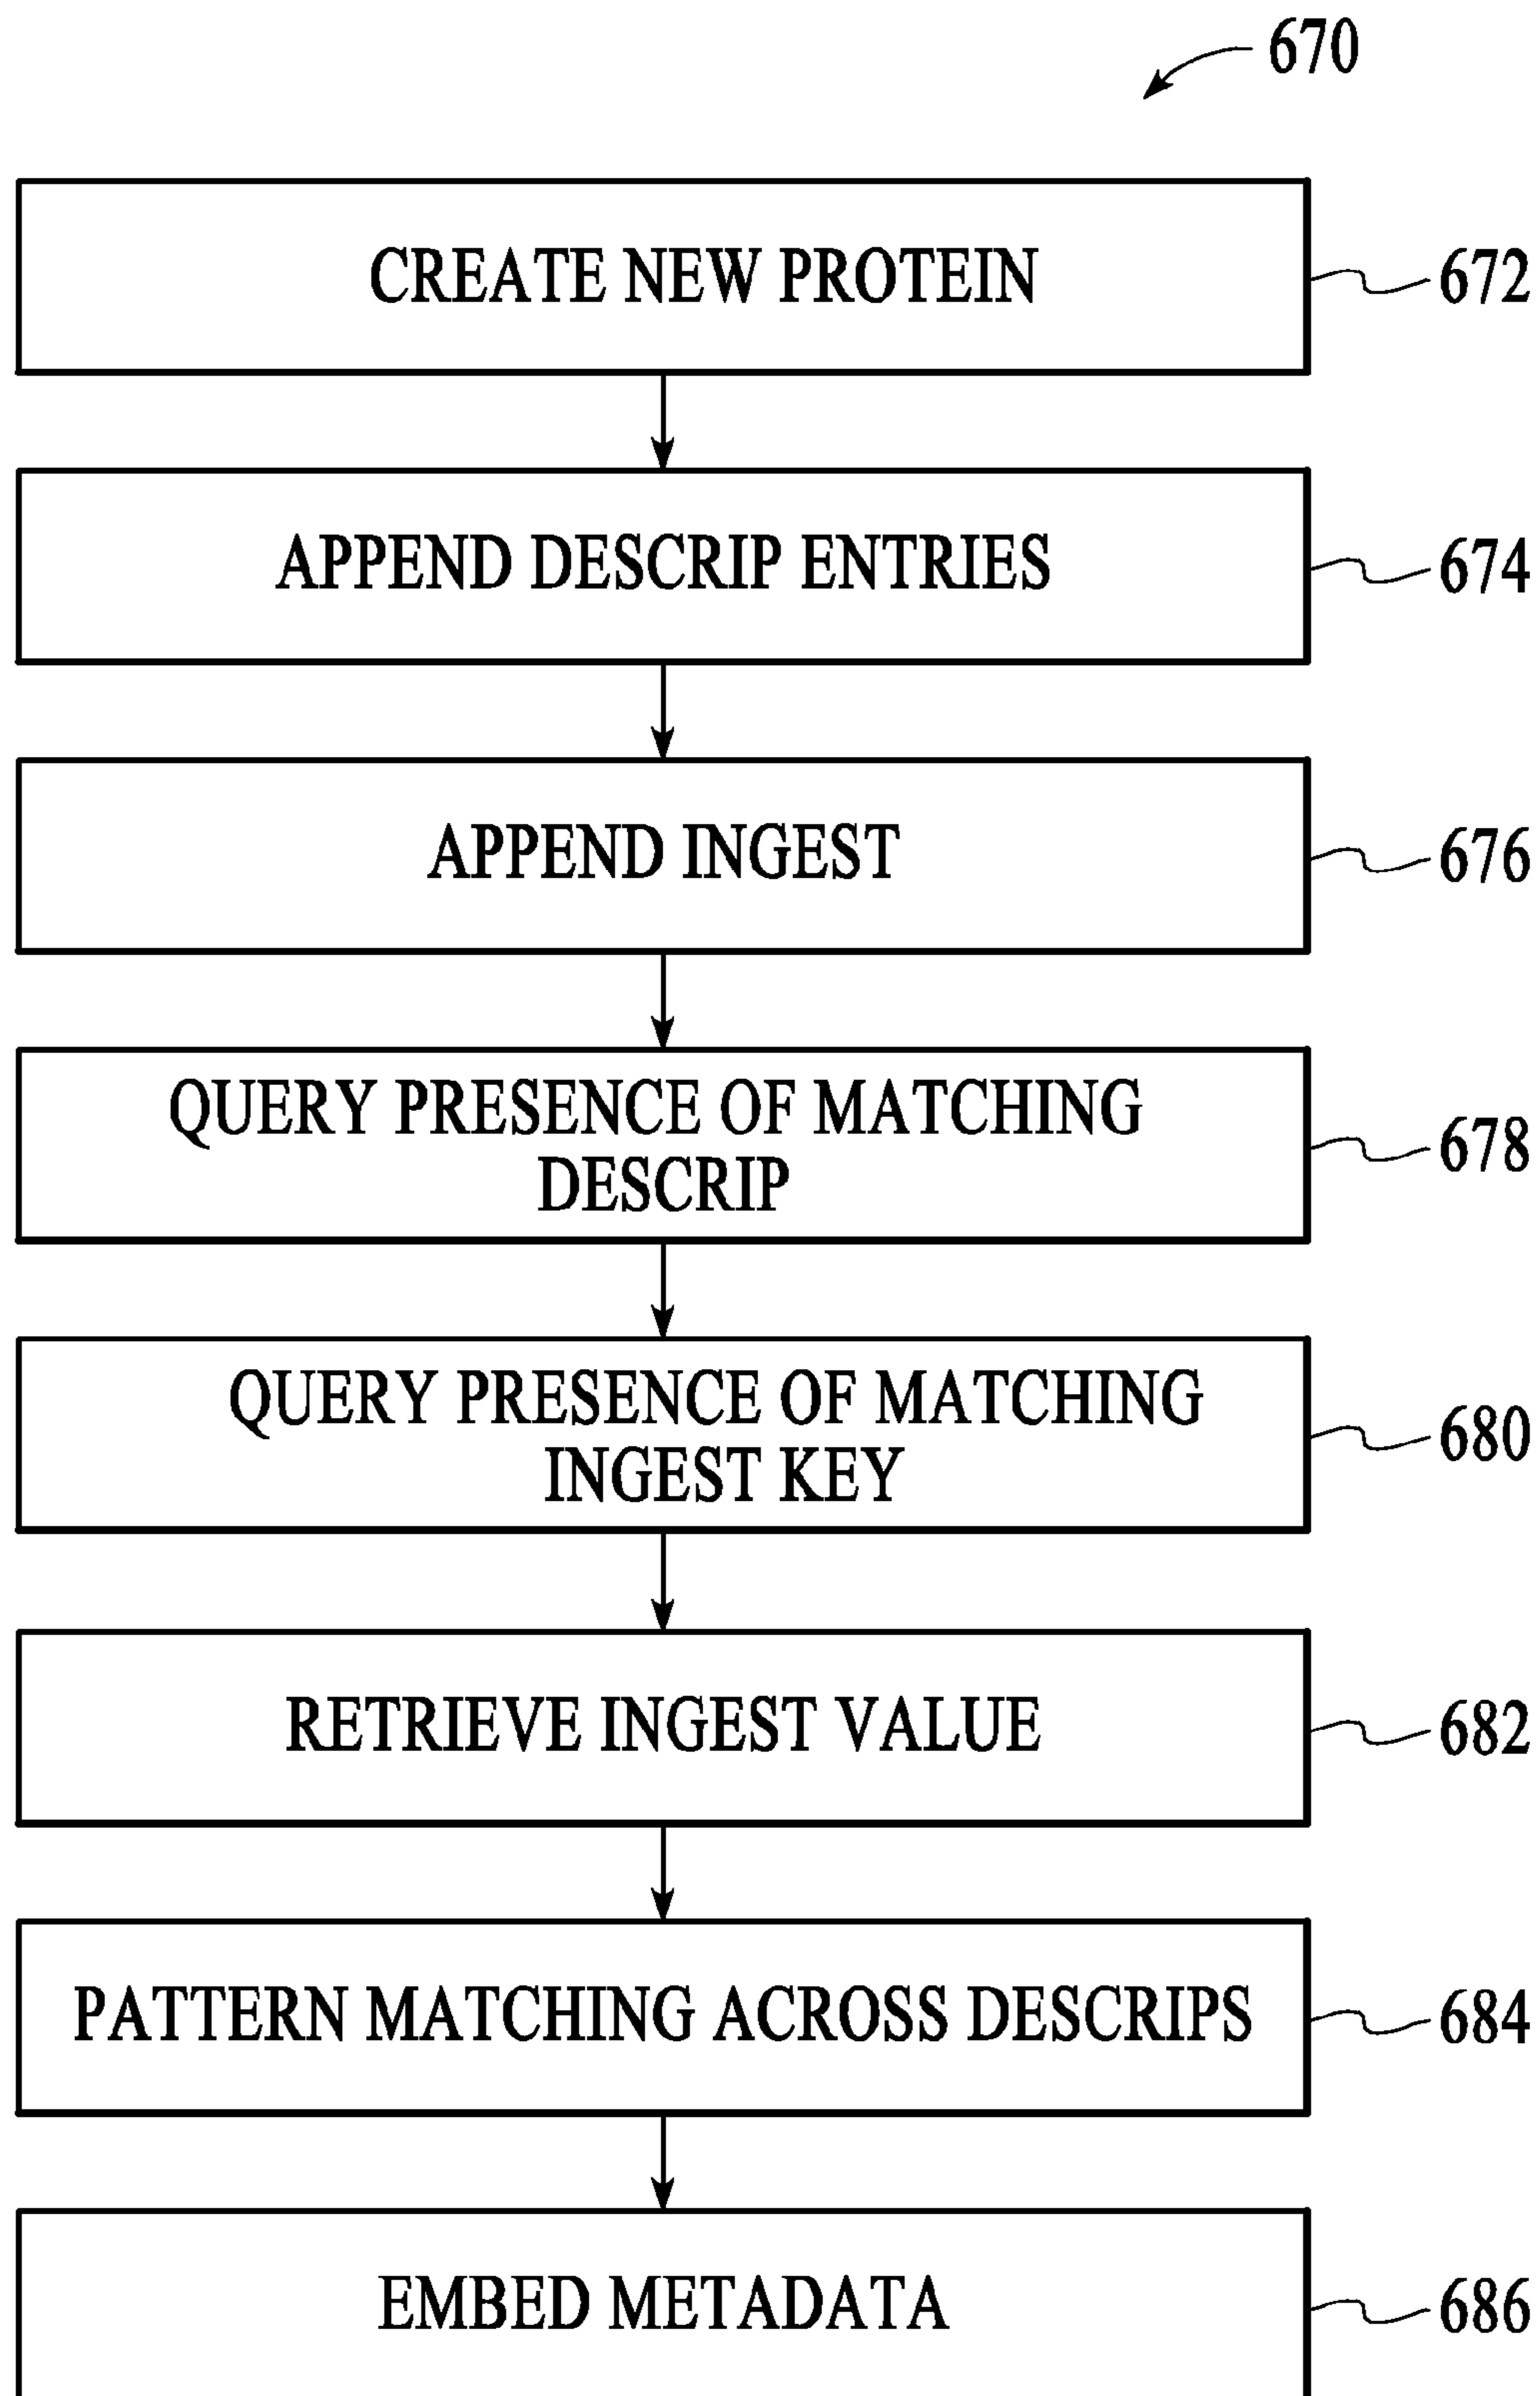


FIG.19D

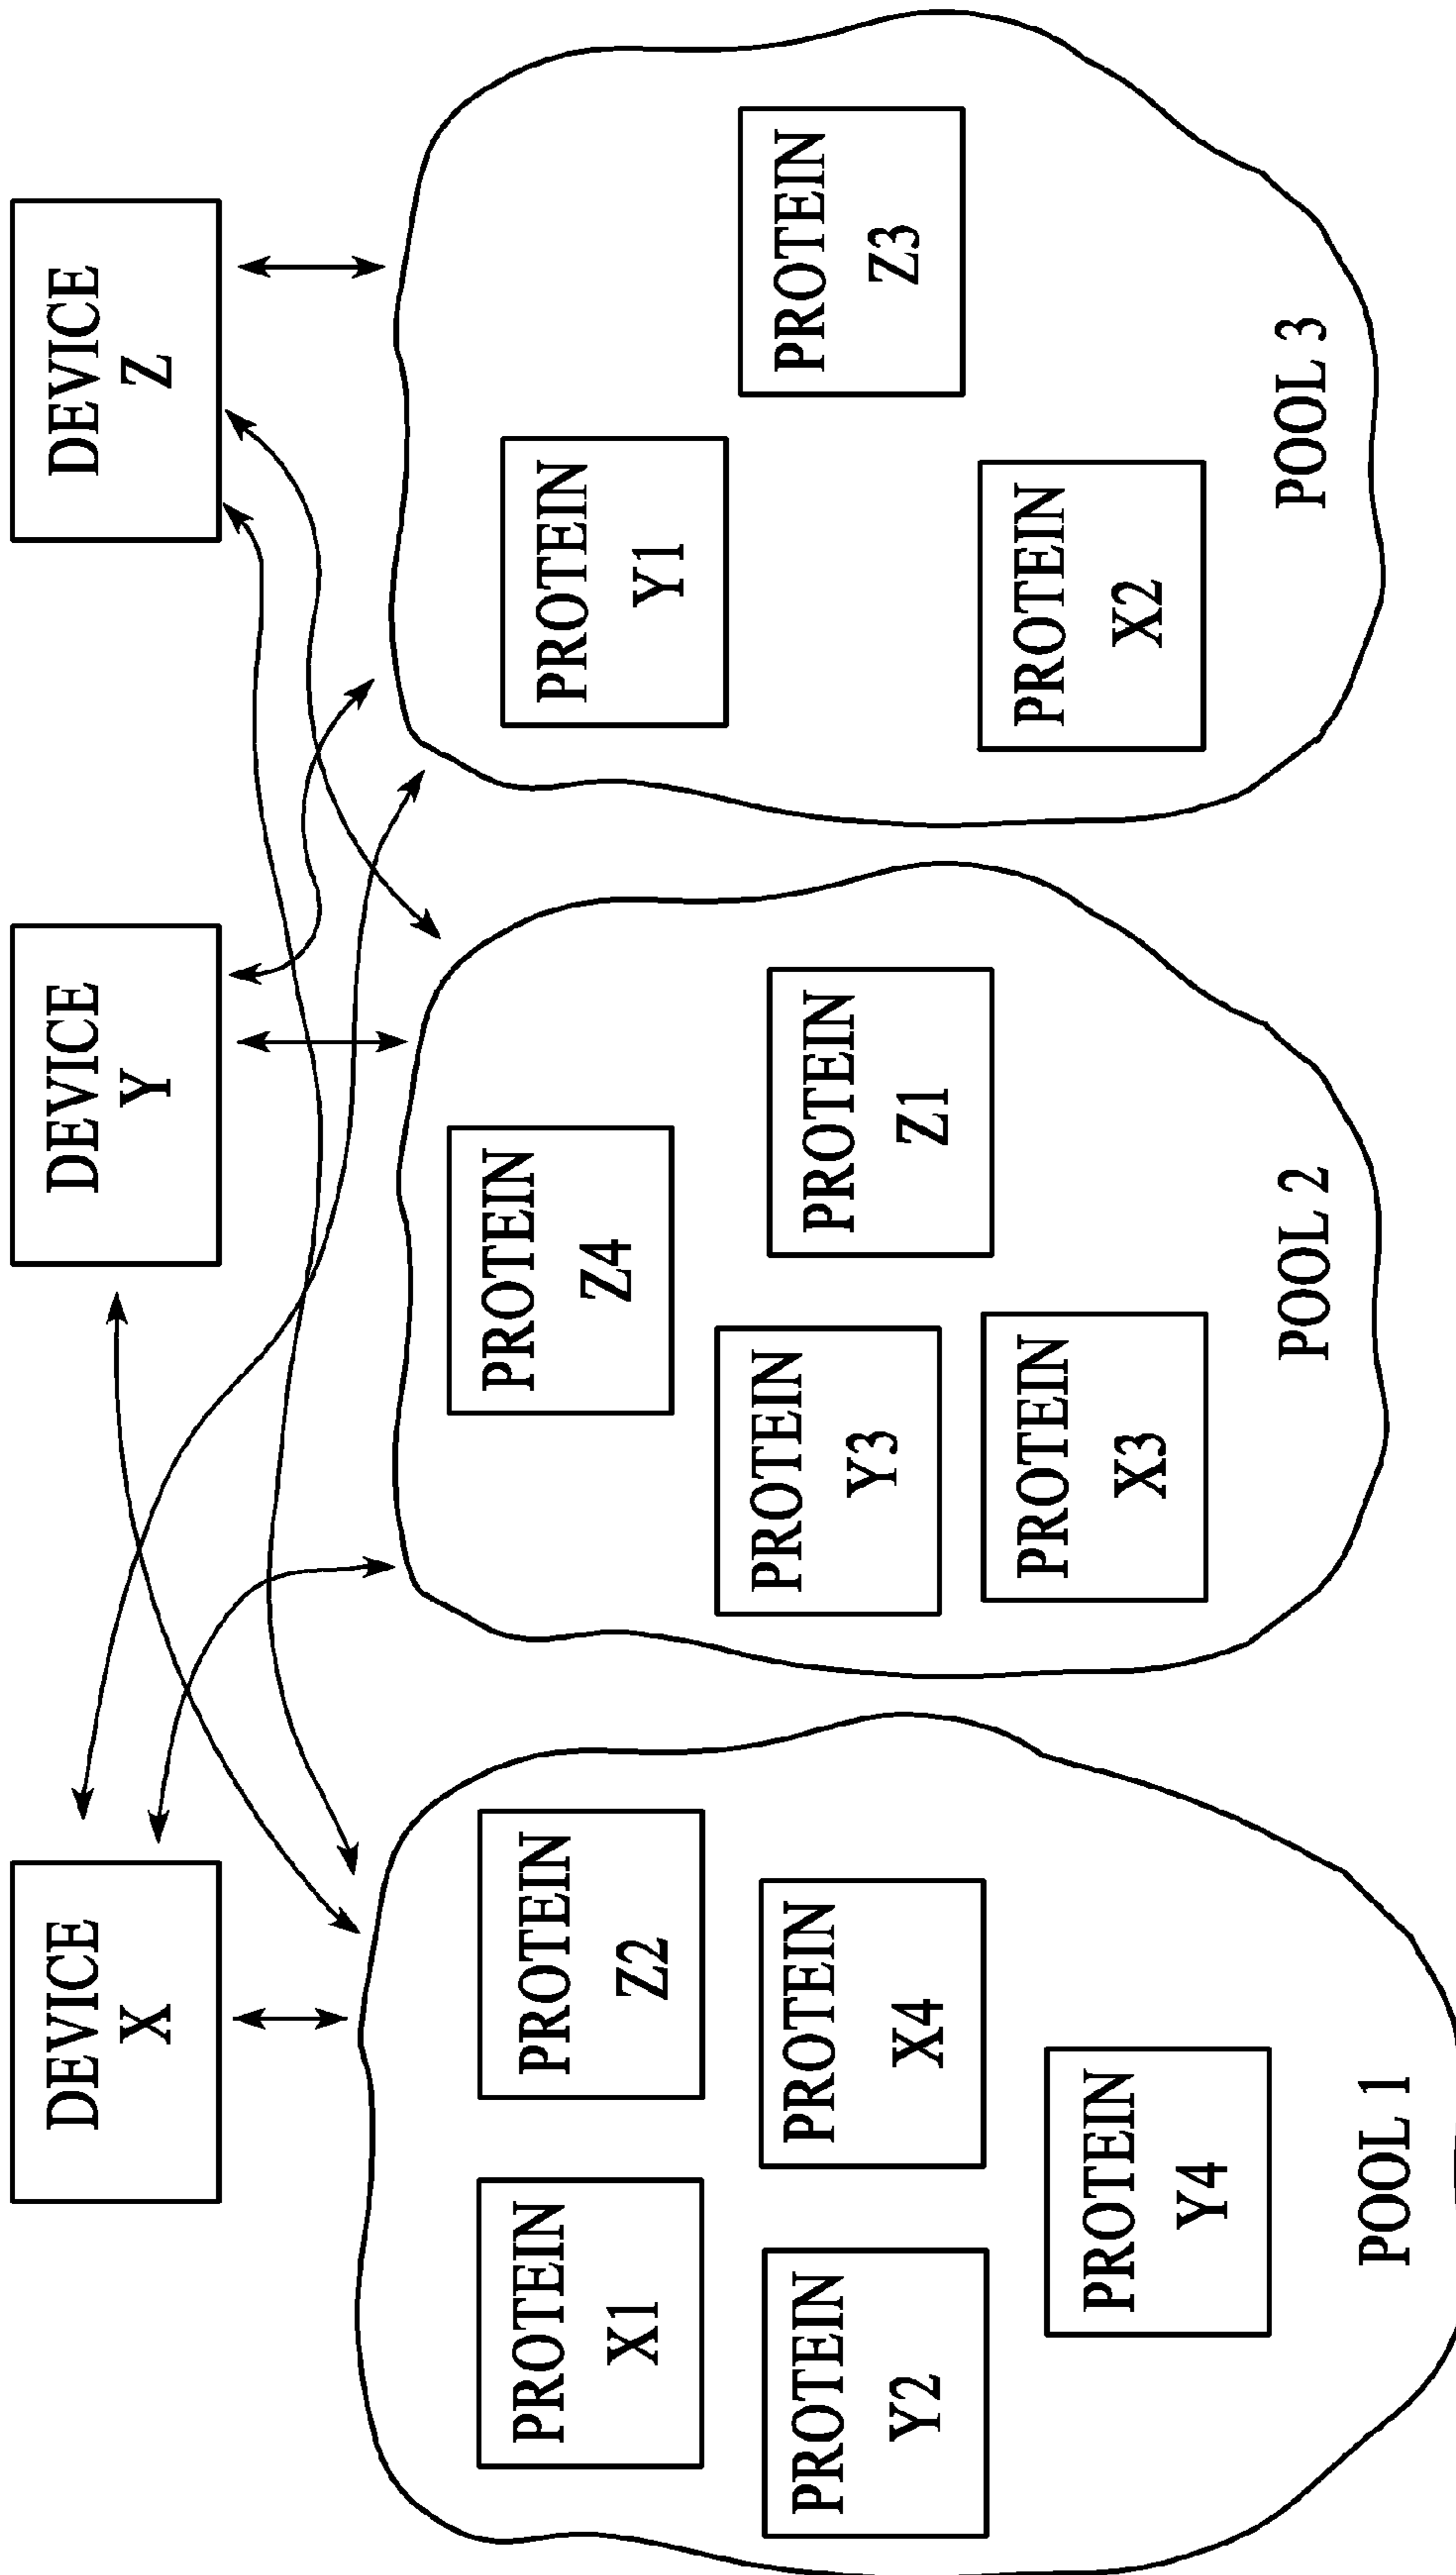


FIG.20

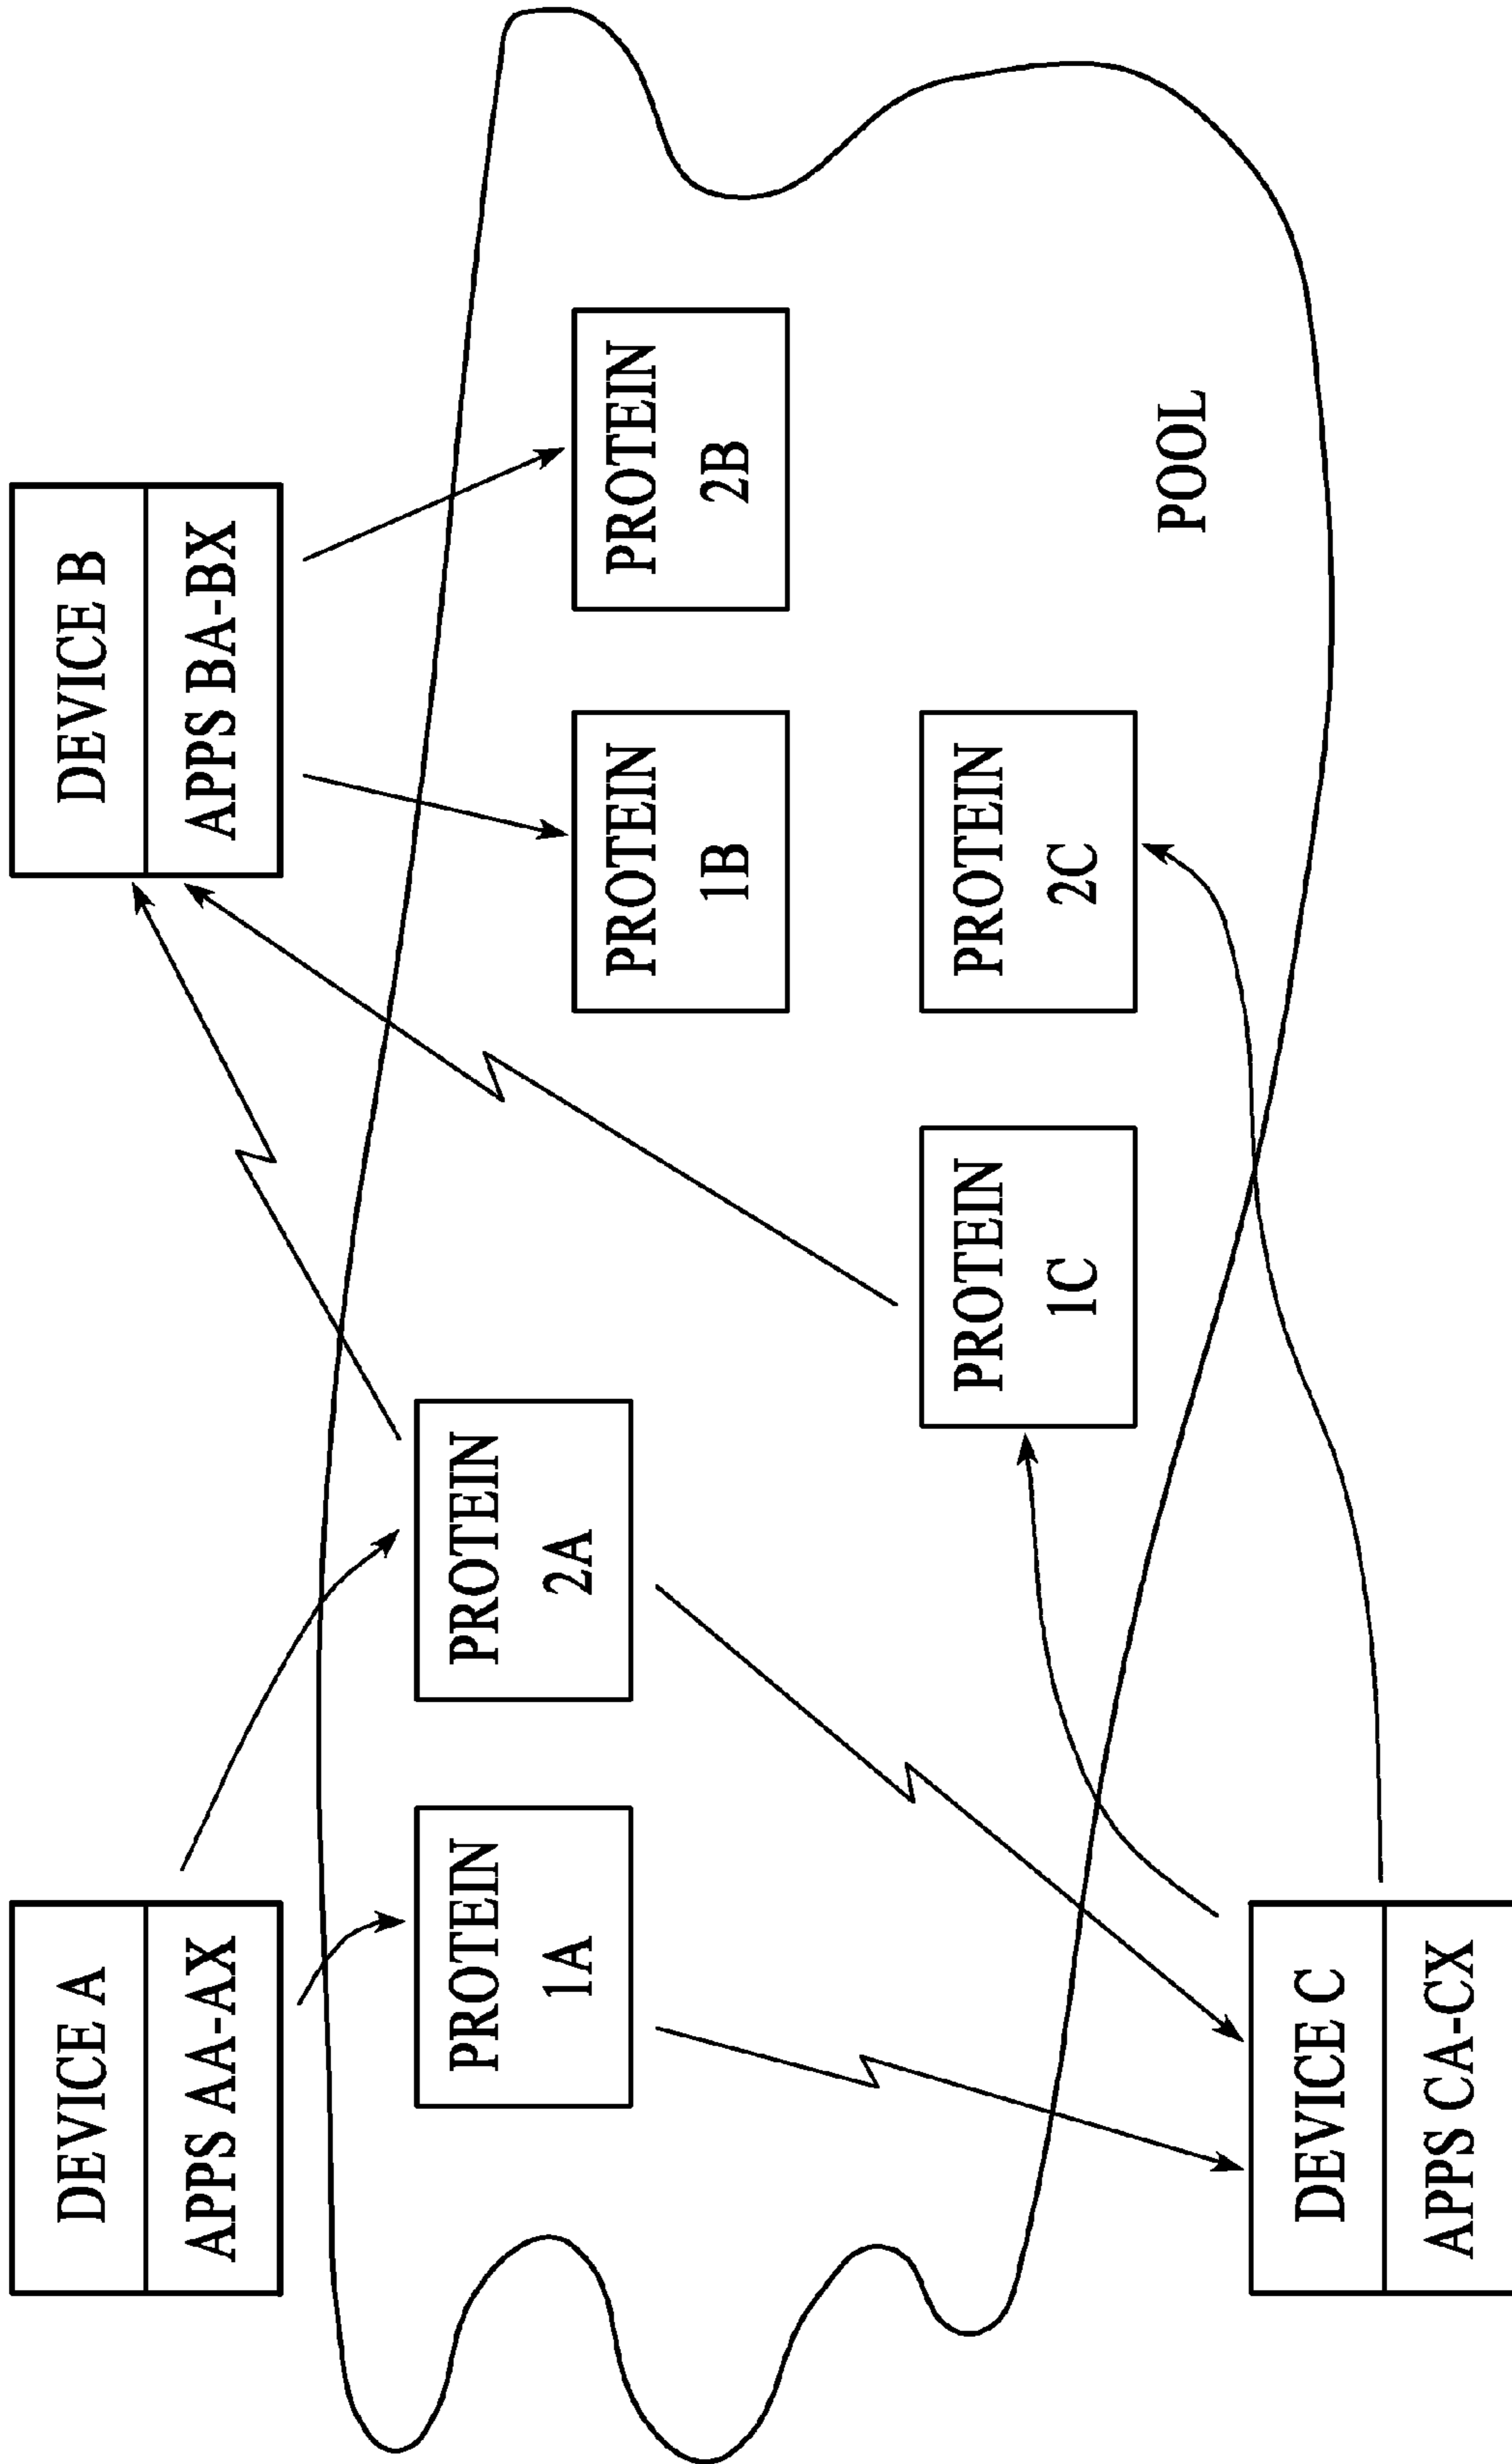


FIG.21

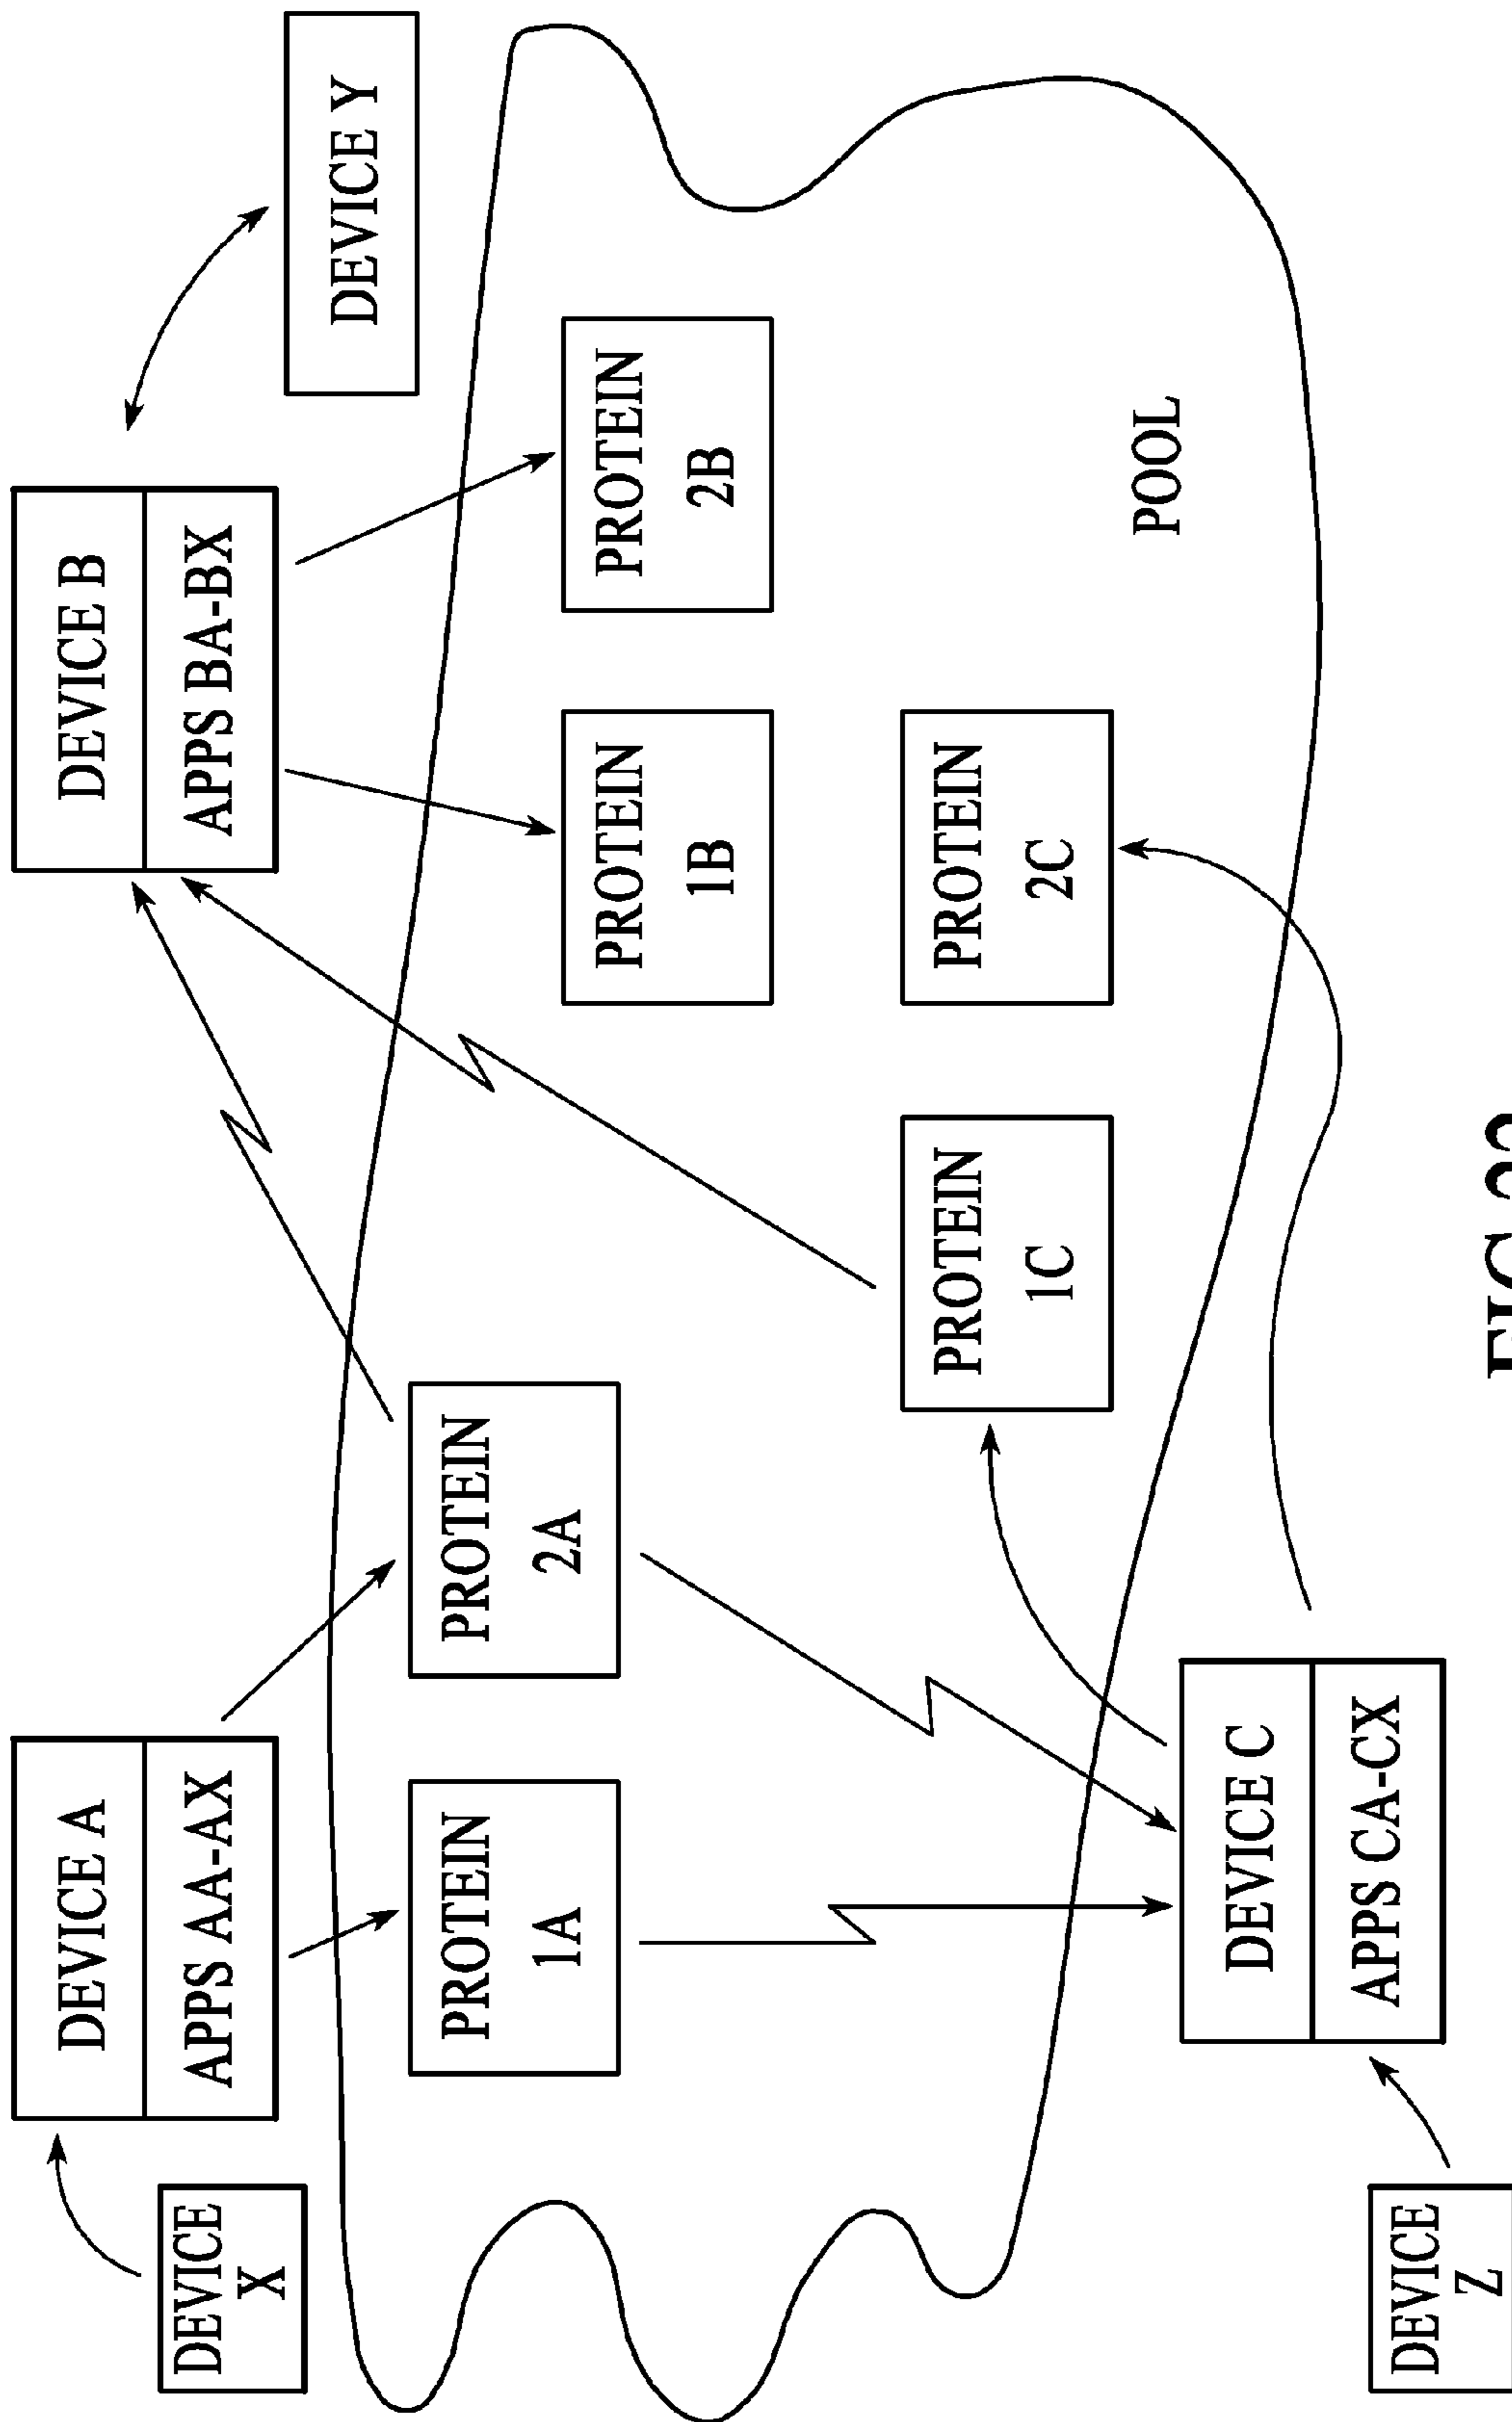


FIG.22

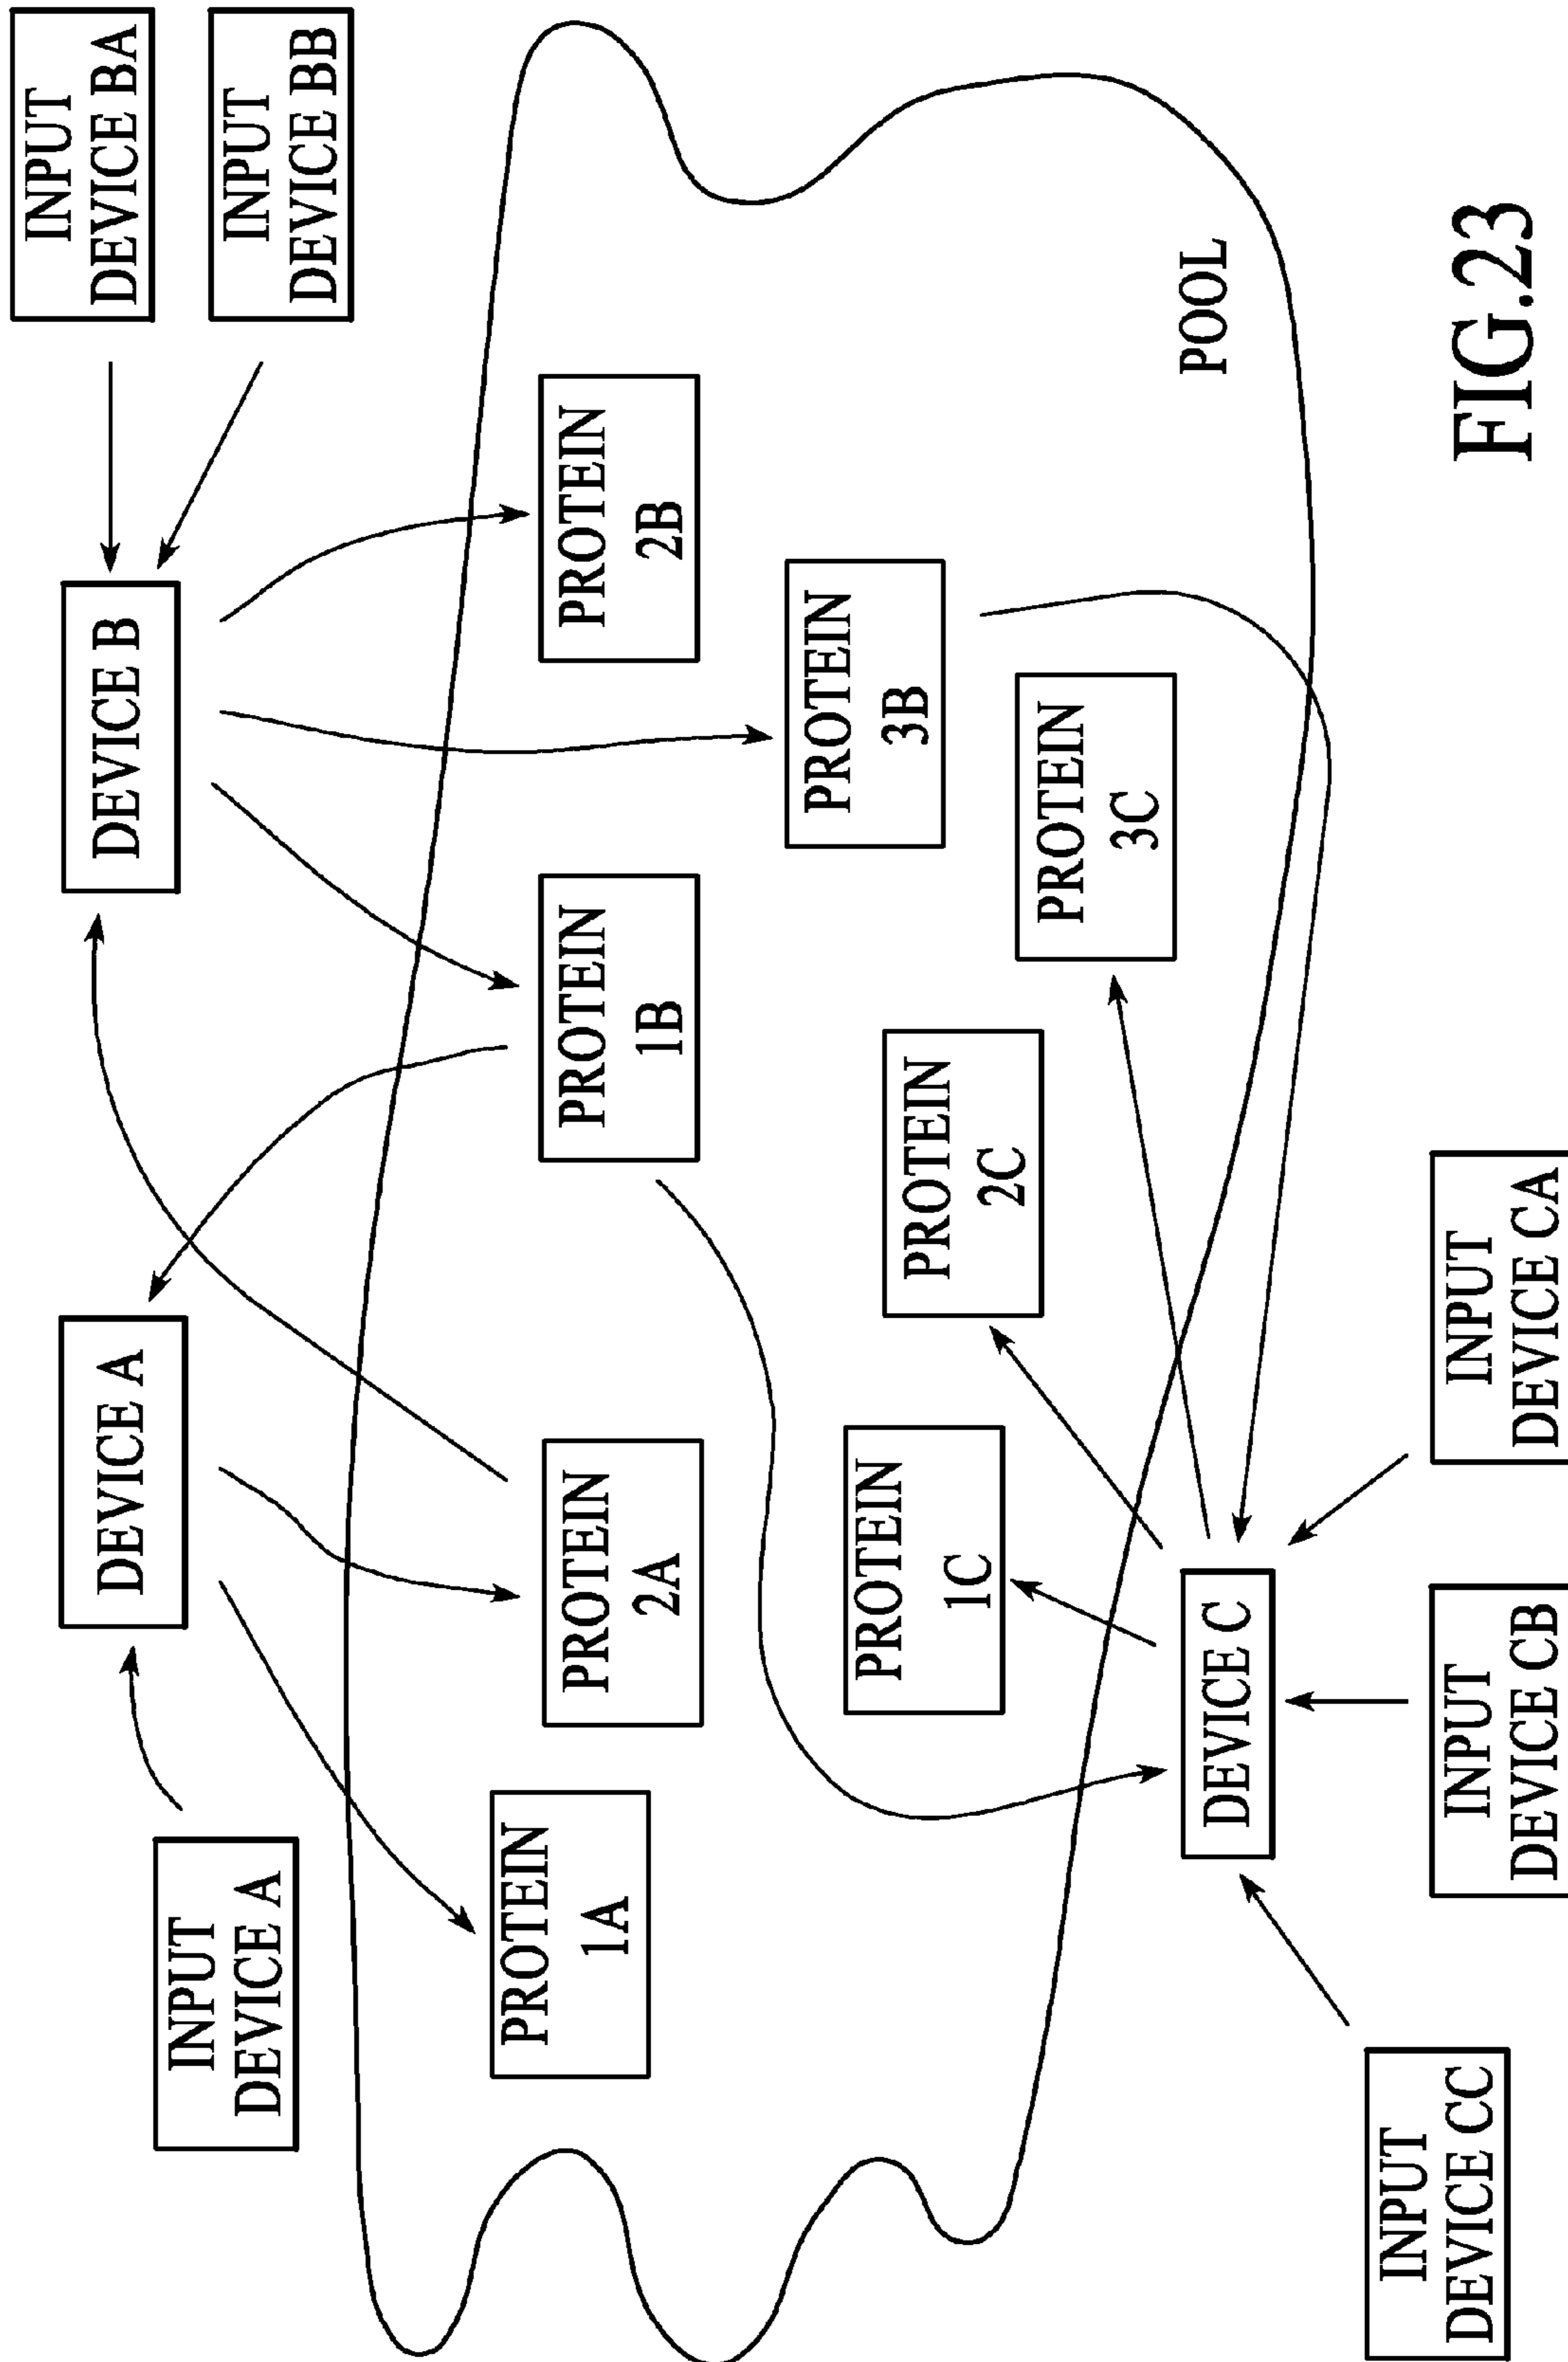


FIG. 23

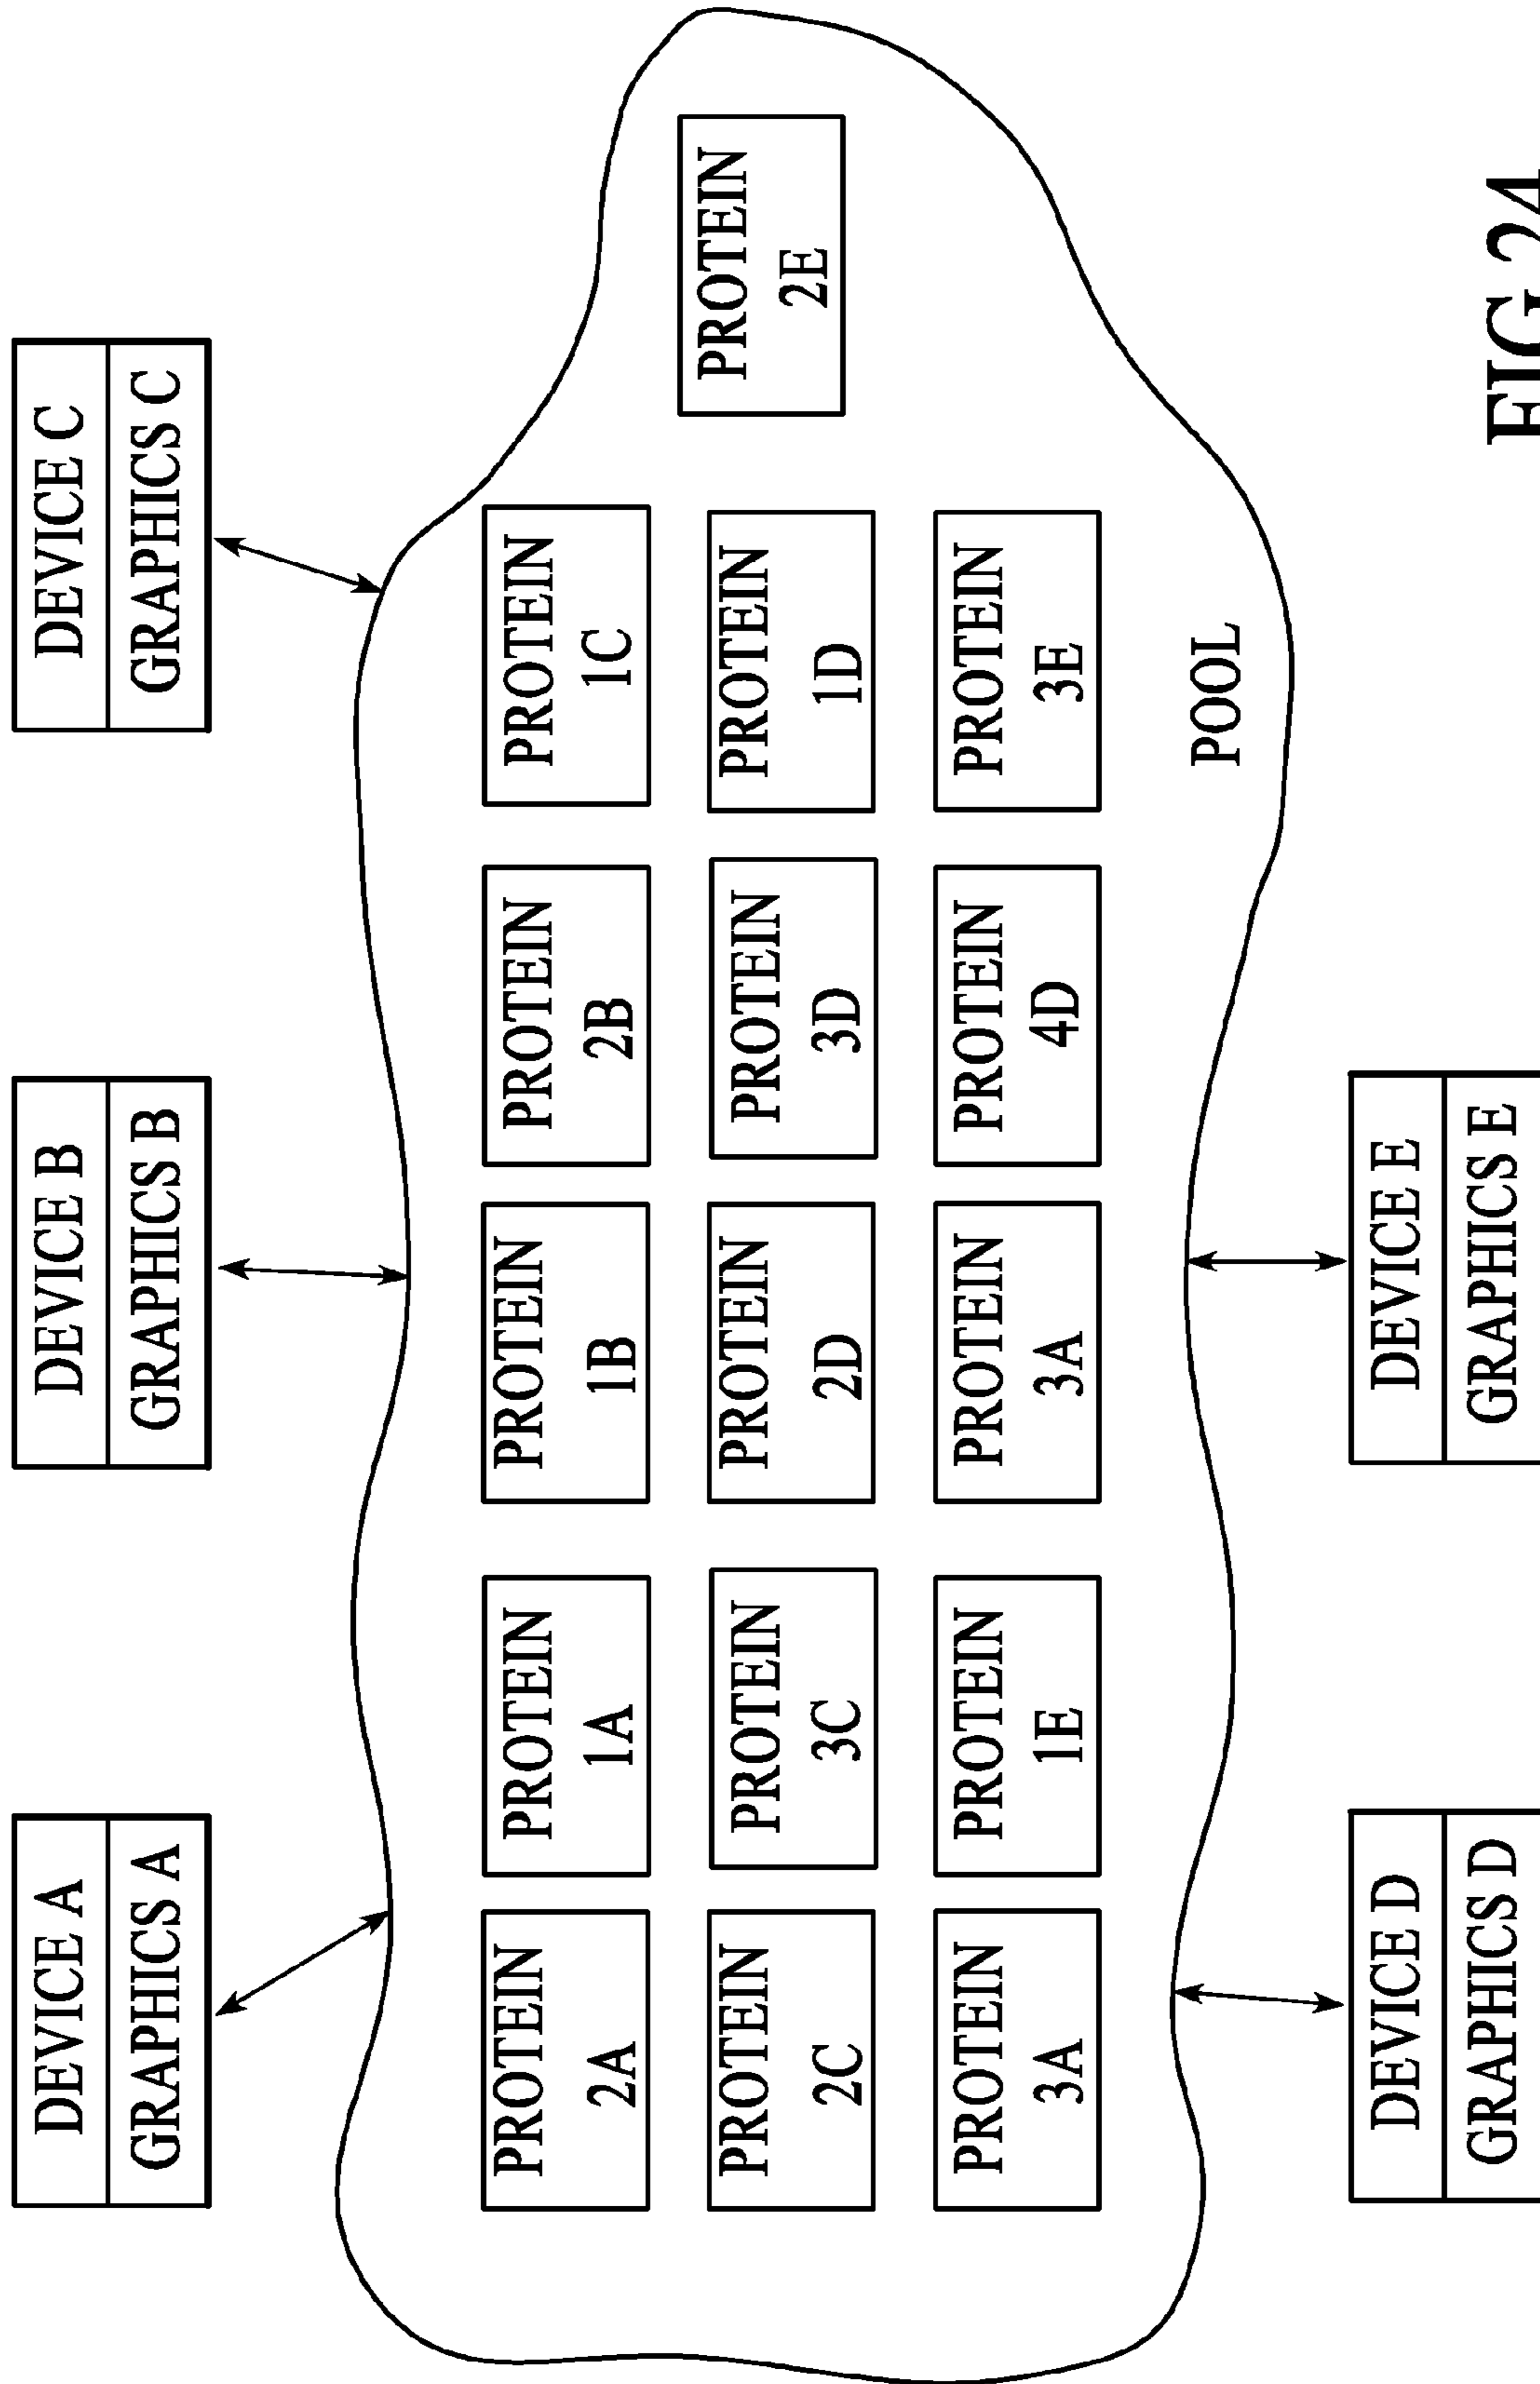


FIG. 24

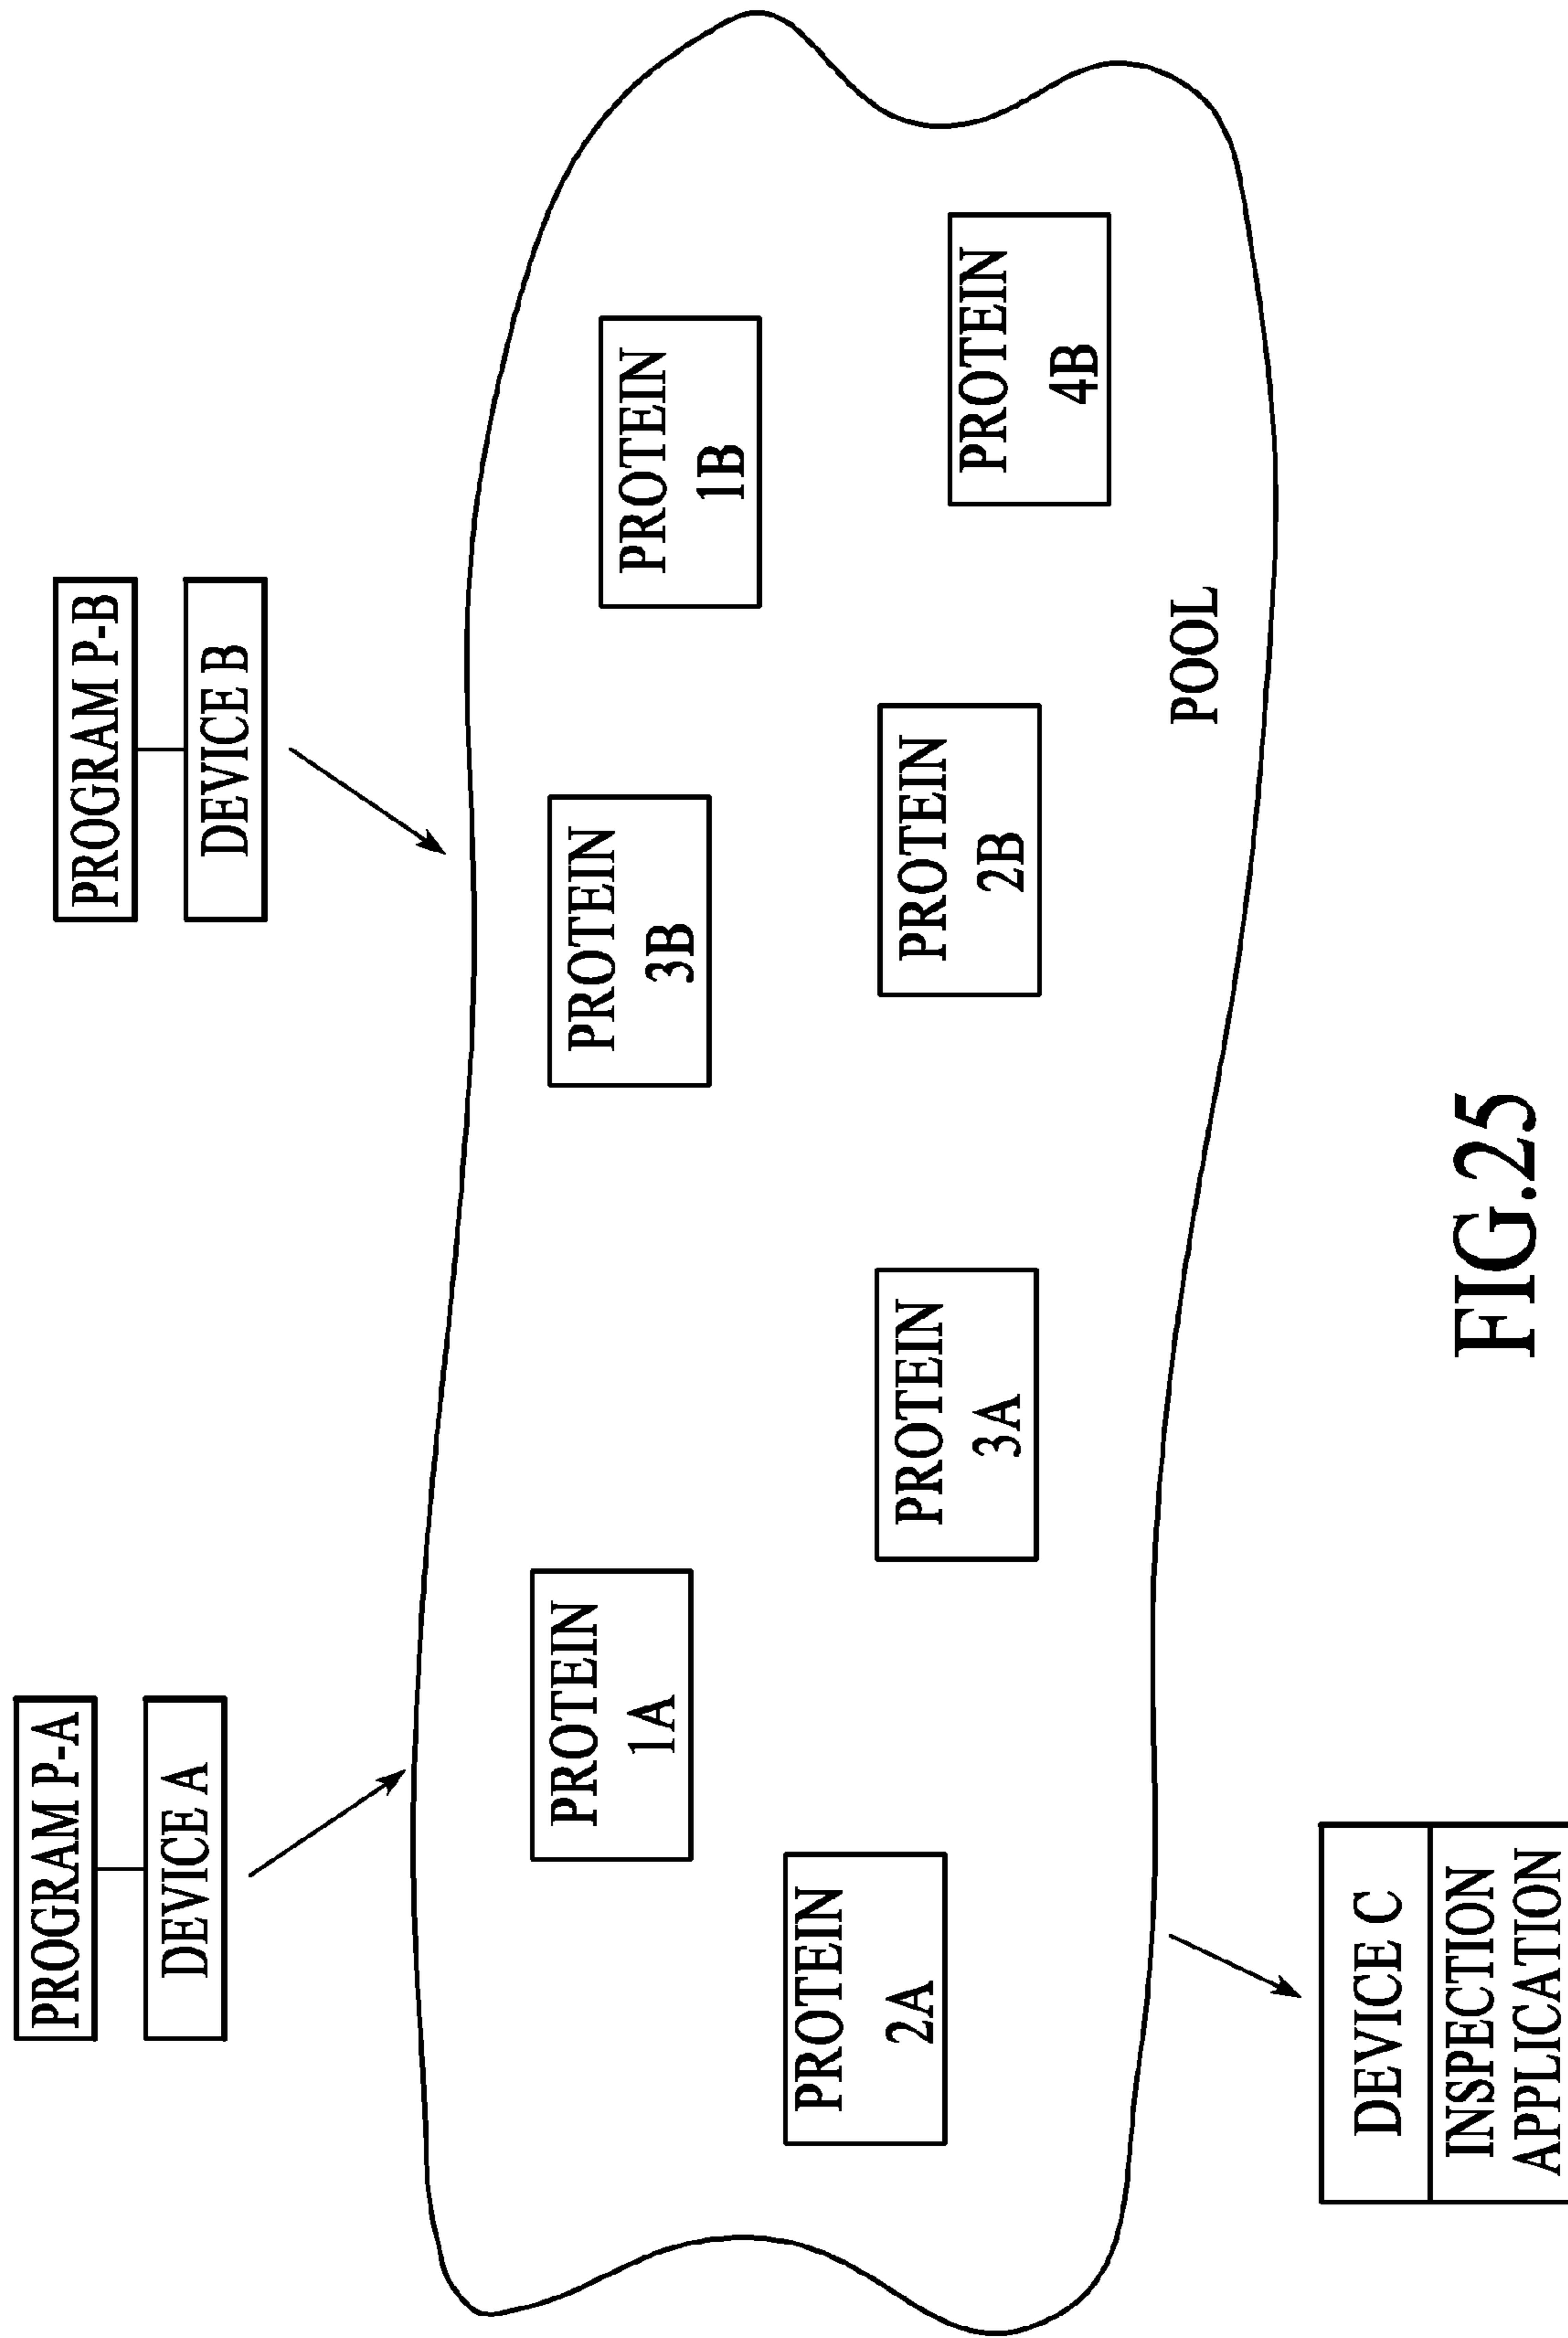


FIG.25

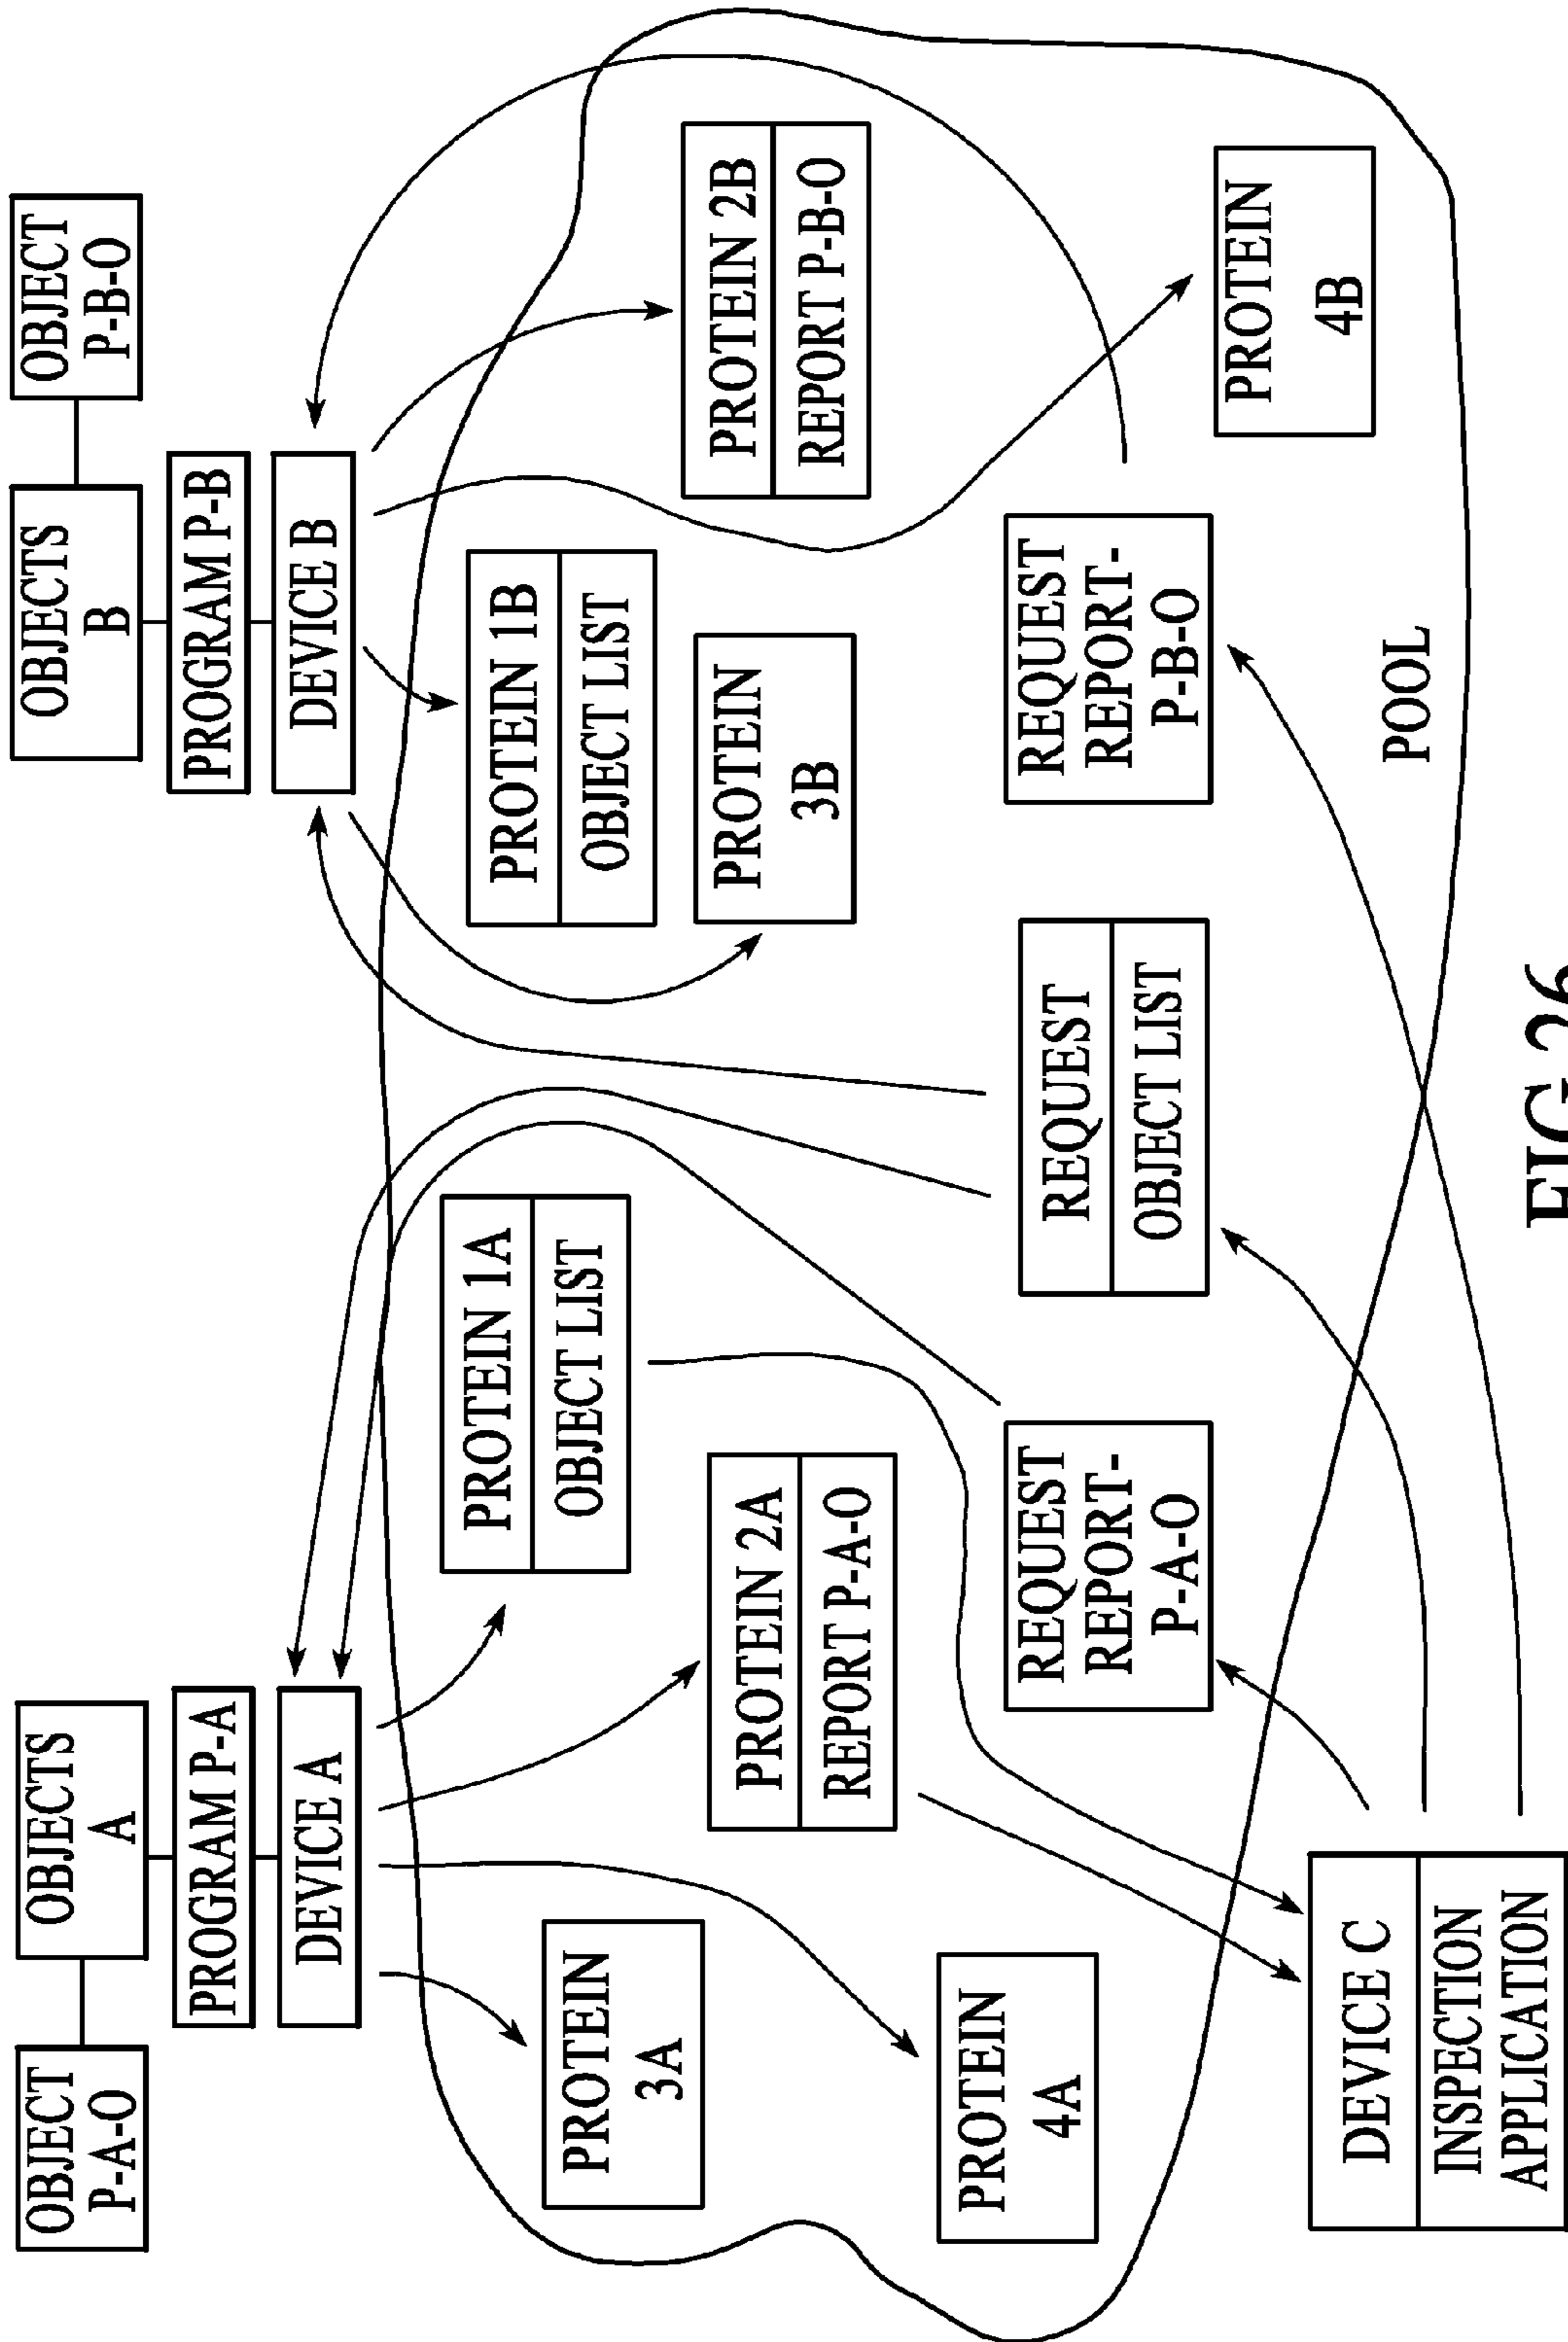


FIG.26

**SPATIAL, MULTI-MODAL CONTROL
DEVICE FOR USE WITH SPATIAL
OPERATING SYSTEM**

RELATED APPLICATIONS

This application claims the benefit of U.S. patent application Ser. No. 61/181,621, filed May 27, 2009.

This application is a continuation in part application of U.S. patent application Ser. No. 11/350,697, filed Feb. 8, 2006.

This application is a continuation in part application of U.S. patent application Ser. No. 12/109,263, filed Apr. 24, 2008.

This application is a continuation in part application of U.S. patent application Ser. No. 12/553,845, filed Sep. 3, 2009.

This application is related to U.S. patent application Ser. No. 12/773,605, filed May 4, 2010.

TECHNICAL FIELD

Embodiments are described relating to control systems and devices including the representation, manipulation, and exchange of data within and between computing processes.

BACKGROUND

Real-time control of computational systems requires the physical actions of a user to be translated into input signals. For example, a television remote control generates specific signals in response to button presses, a computer keyboard generates signals in response to key presses, and a mouse generates signals representing two-axis movement and button presses. In a spatial or gestural input system, the movement of hands and objects in three-dimensional space is translated as signals capable of representing up to six degrees of spatial freedom and a large number of modalities or poses.

INCORPORATION BY REFERENCE

Each patent, patent application, and/or publication mentioned in this specification is herein incorporated by reference in its entirety to the same extent as if each individual patent, patent application, and/or publication was specifically and individually indicated to be incorporated by reference.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows the wand-shaped multi-modal input device (MMID), under an embodiment.

FIG. 2 is a block diagram of a MMID using magnetic field tracking, under an embodiment.

FIG. 3 is a block diagram of the MMID in a tracking environment, under an embodiment.

FIGS. 4a and 4b show input states of the MMID with infrared (IR) light-emitting diodes (LEDs) (IR LEDs), under an embodiment.

FIGS. 5a and 5b show input states of the MMID with IR LEDs, under an alternative embodiment.

FIG. 6 is a block diagram of a gestural control system, under an embodiment.

FIG. 7 is a diagram of marking tags, under an embodiment.

FIG. 8 is a diagram of poses in a gesture vocabulary, under an embodiment.

FIG. 9 is a diagram of orientation in a gesture vocabulary, under an embodiment.

FIG. 10 is a diagram of two hand combinations in a gesture vocabulary, under an embodiment.

FIG. 11 is a diagram of orientation blends in a gesture vocabulary, under an embodiment.

FIG. 12 is a flow diagram of system operation, under an embodiment.

FIGS. 13/1 and 13/2 show example commands, under an embodiment.

FIG. 14 is a block diagram of a processing environment including data representations using slawx, proteins, and pools, under an embodiment.

FIG. 15 is a block diagram of a protein, under an embodiment.

FIG. 16 is a block diagram of a descrip, under an embodiment.

FIG. 17 is a block diagram of an ingest, under an embodiment.

FIG. 18 is a block diagram of a slaw, under an embodiment.

FIG. 19A is a block diagram of a protein in a pool, under an embodiment.

FIGS. 19B1 and 19B2 show a slaw header format, under an embodiment.

FIG. 19C is a flow diagram for using proteins, under an embodiment.

FIG. 19D is a flow diagram for constructing or generating proteins, under an embodiment.

FIG. 20 is a block diagram of a processing environment including data exchange using slawx, proteins, and pools, under an embodiment.

FIG. 21 is a block diagram of a processing environment including multiple devices and numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the devices, under an embodiment.

FIG. 22 is a block diagram of a processing environment including multiple devices and numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the devices, under an alternative embodiment.

FIG. 23 is a block diagram of a processing environment including multiple input devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the input devices, under another alternative embodiment.

FIG. 24 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the graphics events generated by the devices, under yet another alternative embodiment.

FIG. 25 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to allow stateful inspection, visualization, and debugging of the running programs, under still another alternative embodiment.

FIG. 26 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (i.e., pools, proteins, and slaw) are used to

allow influence or control the characteristics of state information produced and placed in that process pool, under an additional alternative embodiment.

DETAILED DESCRIPTION

Systems and methods are described herein for providing multi-modal input to a spatial or gestural computing system. Embodiments of the systems and methods are provided in the context of a Spatial Operating Environment (SOE), described in detail below. The SOE, which includes a gestural control system, or gesture-based control system, can alternatively be referred to as a Spatial User Interface (SUI) or a Spatial Interface (SI).

Numerous embodiments of a multi-modal input device (MMID) are described herein, where the MMID allows the user of a spatial or gestural input system to access a range of input functionalities intuitively and in an ergonomically efficient manner. The MMID of an embodiment is a hand-held input device. The MMID of an embodiment comprises a means of accurately, and in real time, tracking the position and orientation of the device. The MMID of an embodiment comprises a physical and mechanical structure such that the person holding and operating the device may easily rotate it about one or more of its axes. The MMID of an embodiment comprises a physical and mechanical structure such that the device may be held and operated comfortably in more than one rotational grip. The MMID of an embodiment comprises a software component(s) or mechanism capable of interpreting and translating into user input signals both the rotational grip state in which the user is maintaining and operating the device and transitions between these operational rotation states. This software component relies on the tracking data corresponding to the device. In addition, such an input device may have other input capabilities integrated into its form, such as buttons, joysticks, sliders and wheels. The device may also have integrated output capabilities, such as lights, audio speakers, raster displays, and vibrating motors.

As suggested herein, a large variety of specific configurations are possible for the multi-modal input device of the various embodiments. Devices may differ in physical shape, mechanicals, and ergonomics. Devices may also differ in the number of discreet modalities supported by the combination of physical design, tracking technology, and software processing. Furthermore, MMIDs may differ in the design of supplementary on-board input (i.e. beyond position, orientation, and modality), and in on-board output capabilities.

The MMID of an embodiment includes a wand-shaped device with a housing having a form factor similar to a consumer electronics remote control. FIG. 1 shows the wand-shaped MMID 100, under an embodiment. The MMID 100 is approximately five inches long and one and one-half inches wide with a triangular cross-section, but is not so limited. Each face of the MMID 100 housing includes a single input sensor, which in an embodiment comprises an electro-mechanical button, but alternative embodiments can have a greater or lesser number of buttons, or different types of buttons, on each face. When a user holds the MMID 100 one of the triangular prism's long edges 104 naturally faces downward in the user's hand, resting in the bend of the user's fingers, while the prism's opposite face is oriented upward and sits under the user's thumb. The MMID 100 may be rotated 120 degrees about the long axis with a minimal movement of the fingers and thumb, bringing an adjacent face of the prism into the upward orientation. The prism thus includes three distinct, easily accessed modal orientations corresponding to the faces of the prism. The MMID 100 can

be rotated through all (e.g., three) orientations rapidly, repeatably and repeatedly, even by users experimenting with the device for the first time.

Position of the MMID 100 of an embodiment is tracked using magnetic field tracking, as described below, but can be tracked using other tracking technologies (some of which are described herein). The MMID 100 comprises circuitry, a microcontroller, and program code for tracking the device relative to an alternating current (AC) magnetic field, or electromagnetic field (EMF). The EMF of an embodiment is generated or emitted by a compatible base station proximate to the MMID, but is not so limited. The MMID 100 comprises one or more mechanical buttons, also referred to as input sensors, along with corresponding electronics to digitize the state of the one or more buttons. Furthermore, the MMID 100 includes circuitry that provides a radio link to report the tracking data (e.g., orientation data, position data, etc.) and button press raw data to a host system. Additionally, the MMID 100 includes a battery and power supply circuitry.

Input processing software translates the raw tracking and button press data into data comprising six degrees of spatial position and orientation, button down transition, button up transition, and a running account of button state. The input processing software of an embodiment executes in part on the device and in part as application code on the host system, but is not so limited and can run in a distributed manner on any number/combination of processing devices or solely on a single processor. This data is delivered to application software as a series of programmatic "events" (processing of the programmatic events is described in detail below). In addition, this input processing layer provides mode transition and running mode state events to application software. Three states (e.g., i, ii, and iii), and six transitions (e.g., i->ii, i->iii, ii->iii, iii->i, and iii->ii) are possible, as described in detail below.

The processing layer of an embodiment uses hysteresis to allow a user to access a maximum of rotation along the MMID's long axis without leaving a given mode, and to avoid rapid, undesirable flip-flopping between modal states when the MMID is near the edge of a transition angle. Using this hysteresis, to trigger a transition between modes, the MMID of an embodiment should be rotated more than 120 degrees relative to the center angle of the previous mode. So if the MMID is in mode (i), with an absolute angular center of zero degrees, the MMID remains logically in the mode (i) state until a rotation is detected about the long axis of more than, say, 150 degrees in either direction. When the MMID is rotated 151 degrees, it transitions to modal state (ii), which has an angular center of 120 degrees. To effect a return to state (i) the MMID must be rotated in the opposite sense past this angular center by -150 degrees, bringing it past an absolute angle of -30 (or 330) degrees. The hysteresis band, given above as 30 degrees (150 degrees minus 120), is programmatically settable, and may be adjusted by application code or by user preference setting. This hysteresis example is provided for a three-sided MMID, as described above, but is not limited to the values described herein for the three-sided device; the rotation angles and/or hysteresis bands of alternative embodiments are determined according to a form-factor of the housing or wand and to designer/user preferences.

In addition, certain modes can be selectively disabled by application code. So the MMID can be treated by application code as a single-mode device outputting a constant modal state of (i), (ii), or (iii). Or, any one of the modes may be disabled, either by mapping the disabled mode to either of the two remaining modes exclusively, or by treating the disabled mode as an additional area of the hysteresis band.

Further, the system may be configured to immutably associate a physical face of the MMID (e.g., triangular prism) with each mode, the faces being optionally labeled as to mode association by means of active or passive markings. Alternatively, the system may be configured to assign modes to faces in a contextual way. As an example of this latter case, the MMID can be configured so that, when it is first picked up by a user after a period of inactivity, the initially upward face is associated with mode (i). In such cases an indicator of the active mode can be provided on the MMID, on the graphical display to which the user is attending, or on a combination of the MMID and the graphical display.

Each face of the MMID includes a single button, also referred to as an input sensor. These buttons are treated identically by application-level software, but are not so limited. From the user's perspective, the device may be considered as having a single logical button, with three physical incarnations for reasons of ergonomic practicality. The circuitry and software of the MMID does distinguish manipulation of different physical buttons, however, and the system may be arranged so that pressing the buttons in specific combinations places the device in various configuration and reset states.

The MMID of an embodiment functions using magnetic field tracking technology (see, for example, U.S. Pat. No. 3,983,474). The use of orthogonal coils for generating and sensing magnetic fields has been used in locating and tracking remote objects. For example, U.S. Pat. No. 3,644,825 teaches generating and sensing coils which move with respect to each other. Alternatively, the magnetic field can be made to rotate as taught in Kalmus, "A New Guiding and Tracking System", IRE Transactions on Aerospace and Navigational Electronics, March 1962, pages 7 through 10.

The use of coordinate transformers to determine the orientation of a first coordinate system with respect to a second coordinate system has also been used. For example, U.S. Pat. Nos. 3,474,241 and 3,660,648 disclose transformers which transform angular rates or angular errors measured in a first coordinate frame into angular rates defined about the axes of an intermediate coordinate frame about whose axes the angular rotations or rates are defined and then integrate to determine the angles defining the angle-axis sequence which defines the orientation of the first coordinate frame with respect to a second coordinate frame through the use of Euler angles.

FIG. 2 is a block diagram of a MMID using magnetic field tracking, under an embodiment. A base station 210 located proximate or in the tracking environment of the MMID both provides the tracking field, as well as communicates with the MMID 211. In the base station, a signal generator creates magnetic fields by using a field generator circuit 201 to produce a wave form alternately in three orthogonal coils 202. The electromagnetic signals generated by these coils are received by three orthogonal coils 203 in the MMID. The received signals from the three coils are typically amplified using operational amplifiers 204 and converted to digital signals 205 which can be sampled by a microprocessor 207. The microprocessor analyzes the input of the three coils using digital signal processing (DSP) techniques. The DSP process provides a location vector projecting the distance and direction of the MMID from the base station, as well as an orientation matrix that determines the orientation of the MMID.

Additional information (e.g., time stamp, universal ID, etc.) can also be combined with the MMID location data. One or more user input sensors 206 are also sensed for state. The input sensors 206 can be momentary switches, toggle switches, joystick style input devices, and/or touch sensors to name a few. The sample data from these switches includes a

single bit (for a touch button) or a more complex data value, such as a floating point x,y coordinate for a touch sensor.

In an embodiment, the microprocessor communicates data including location data and orientation data from the MMID wirelessly to a host process. The MMID has a radio frequency transmitter and receiver (TX/RX) 208 for data communication to the network through an Access Point 209. This radio link can use any wireless protocol (e.g., Bluetooth, 802.11, Wireless USB, proprietary solutions, Nordic Semiconductor nRF24L01 low power radio solution, etc.). The access point can communicate the received data stream to one or more host computers through a local area network (e.g., Wired Internet 10/100/1000 BaseT, 802.11, etc.) or other interface (e.g., USB, etc.).

FIG. 3 is a block diagram of the MMID in a tracking environment, under an embodiment. The MMID 304 is shown in relation to the tracking environment 300. The MMID is communicating with a base station 301, as described above, but the MMID can communicate with any number of different types and/or combinations of electronic devices in the tracking environment 300. The tracking environment is not limited to a particular size because, as the range of the radio frequency communications channel may be different from the range of the AC magnetic field, additional AC magnetic field generators 305/306/308 with coils can be provided to create additional tracking beacons. These beacons can operate at different frequencies and/or transmit at different times. As the user of the MMID moves away from field generator 302 and towards generator 305 the MMID will use whichever signal is instantaneously stronger to determine location and orientation, but will still communicate this data back to the network using access point 303.

As the MMID moves out of range of the access point 303 and towards base station 306, the MMID will associate the radio link with the access point in base station 306. The ability to roam among magnetic field generators and data access points ultimately allows the MMID to be used in an arbitrarily large tracking environment. Note that the access points and magnetic field generators need not be at the same location 307/308. While both the access points and field generators have means of communication with one or more host devices over a local area network, the frequency generators can operate autonomously 305 allowing for easier installation.

Following is an operational example of a person using the MMID of an embodiment. During operation, an operator stands some distance (e.g., ten feet) before a triptych-format wide aspect ratio projection screen, roughly two meters high and four meters wide; a one-point-five meter wide table stands immediately before her. The table is itself also a projection surface treated by a projector ceiling-mounted immediately overhead. The operator holds the MMID having the triangular-cross-section MMID comfortably in her right hand, with flat side "i" pointing upward. As she aims the MMID toward and about the front screen, a partially transparent graphical cursor indicates the intersection of the MMID's pointing vector with the screen surface. The input system's high frame rate and low latency contribute to a strong sense of causal immediacy: as the operator changes the MMID's aim, the cursor's corresponding movement on the forward screen does not apparently lag behind; the perception is of waving a flashlight or laser pointer.

The application in use by the operator is a product packaging preview system, and is configured to make use of the MMID in a way identical to many similar applications; the MMID modalities are thus well familiar to the operator. Mode "i" allows direct manipulation of application elements at the fully detailed level; mode "ii" performs meta-manipulation of

elements (e.g. at the group level); and mode “iii” permits three-dimensional manipulations. At any instant, the appearance of the cursor reflects not only the current mode but also indicates visually the direction of axial rotation that would be necessary to switch the MMID’s modes. At present, the cursor shows that a clockwise rotation of the MMID would cause a modal transition to “ii”, while counterclockwise rotation would transition to mode “iii”.

Arranged on the left third of the forward screen triptych is an array of small object groupings. The operator rotates the MMID axially clockwise until the next face is aimed upward, under her thumb, and the cursor changes to indicate the modal transition to state “ii”. She aims the MMID leftward, and as the cursor travels over each object grouping a highlight border fades up, subsequently fading down as the cursor exits the grouping’s convex hull. The operator allows the cursor to rest on a particular grouping and then depresses the button immediately under her thumb. The cursor indicates that the object grouping has been grabbed and, as she swings the MMID toward the center of the forward screen, the grouping moves so as to track along with the cursor. The operator releases the button when she has brought the miniature grouping to a position directly in front of her. The grouping rapidly expands to fill the full extent of the center third of the forward screen, revealing a collection of variously shaped plastic bottles and the textual indication “Pet Energy Beverages”.

The operator once again rotates the MMID clockwise about its long axis, whereupon the cursor changes to indicate that mode “iii” is now operational and, thus, that 3D manipulation is enabled. The operator aims the cursor at a particularly bulbous bottle shaped like a coiffured poodle leg, and the bottle visually highlights; the operator then depresses the button. The system now enters a direct-manipulation mode in which translation and rotation of the MMID controls translation and rotation of the selected object in the virtual space being rendered. So, as the operator pulls the MMID toward herself (directly along the geometric normal to the forward screen), the bottle grows larger, verging toward the virtual camera. Similarly, left-right movement of the MMID translates to left-right movement of the rendered bottle (along the screen’s lateral axis), and up-down translation of the MMID results in vertical translation of the bottle. An appropriate scale factor, customizable for each operator, is applied to these translations so that modest movements of the MMID effect larger movements of virtual objects; the full extent of the graphical/virtual environment is thereby made accessible without exceeding an operator’s range of comfortable hand-movement.

A similar scaling function is applied to the mapping of MMID orientation to absolute rotational position of the rendered bottle. In the present example, the operator’s preferences dictate a four-times scale, so that a ninety degree rotation of the MMID around any axis results in a full three hundred sixty degree rotation of the virtual object (90 degrees multiplied by four (4) results in 360 degrees). This insures that wrist- and arm-based MMID rotations remain within a comfortable range as the operator examines the bottle from every possible angular vantage. So, for example, as she rotates the MMID upward, tipping it ninety degrees around a local x-axis so that it evolves from forward-pointing to upward-pointing, the bottle executes a full rotation around the screen-local x-axis, returning to its initial orientation just as the MMID achieves a fully upward attitude. Note that an appropriate mode-locking effect is applied so long as the MMID’s button remains depressed: the operator may rotate the MMID one hundred seventy clockwise degrees around

the MMID’s long axis (producing a five hundred ten degree “in-screen” rotation of the virtual object) without causing the MMID to switch to mode “i”.

When the operator releases the MMID’s button, the rendered bottle is released from direct manipulation and retains its instantaneous position and rotation. If at the moment of button release the MMID is in a rotational attitude that would ordinarily correspond to a MMID-mode other than “iii”, the operator is granted a one-second temporal hysteresis (visually indicated as part of the on-screen cursor’s graphical state) before the mode switch is actually effected; if the operator returns the MMID rotationally to an attitude corresponding to mode “iii”, then direct 3D manipulation mode persists. She may then perform additional positional and attitudinal adjustments by superimposing the cursor atop the bulbous bottle and again depressing the button; if instead she aims the cursor at a different bottle, that object will be subject to her manipulations.

The operator eventually switches the MMID to mode “ii” and, using a dragging modality identical to that by which she brought the bottle grouping to the center screen, brings a color-palette from the right screen to the center screen; when she releases the button, the palette expands and positions itself to the side of the bulbous bottle. She then rotates the MMID to select mode “i” and manipulates the color palette’s selection interface; when the crimson hue she desires has been selected, she depresses the button and drags a color swatch from the palette downward and leftward until it overlies the clear material forming the bulbous bottle. When she releases the button, the color is applied and the bottle’s material adopts a transparent crimson.

Still in mode “i”, the operator points the MMID directly at the bulbous bottle, which highlights in response, and, depressing the button, swings the MMID downward to drag the image of the bottle from the front screen to the surface of the table immediately before her. She releases the button and thereby the bottle, leaving it in position on the table. The operator then rotates back to mode “ii” and points the MMID forward at the collection of other pet energy beverage bottles; she depresses the button and immediately flicks the MMID leftward, releasing the button a fraction of a second later. The collection of bottles flies leftward, diminishing in size as it travels, until it comes to rest in the location and at the overall scale at which it started. The operator then selects a different grouping of pet care products, bringing it to the center display region as before in order to select, inspect, and modify one of the items. She eventually adds the selected object to the table display. The operator continues this curatorial process.

At a certain point, the operator elects to modify the physical geometry of a canister of pet massage oil using a simple geometry editor, also pulled from the collection of tools appearing on the right third of the forward screen triptych. The description of many manipulations involved in the use of this editor is omitted here, for the sake of clarity, except as regards the simultaneous use of two MMIDs. In the present instance, the operator uses a second MMID, held in her left hand, in order to put a twist in the canister (originally a simple extruded shape with rectangular cross section) by using one MMID to grab the top part of the canister’s geometry and the other MMID to grab the canister’s bottom part (both MMIDs in mode “iii”). With the top and bottom thereby separately “affixed”, the operator rotates the MMIDs in opposite directions; this introduces a linear twist about the canister’s main axis. The operator finishes these geometry modifications and returns the editing module to the right display; she adds the modified canister to the table’s growing assortment.

At last there are a dozen objects being rendered on the table, and the forward center display is empty once more—the operator has mode-“i”-flicked the last grouping leftward (and the color palette rightward). She then points the MMID, still in mode “i”, at the table, but her aim avoids the product renderings there; instead, she depresses the right button and describes a circular trajectory with the MMID, as if drawing a curved corral shape around the displayed objects. In response, the system applies a grouping operation to the formerly distinct product renderings, organizing their layout and conforming their relative sizes. Finally, the operator uses mode-“ii”-dragging to elastically extend the input aperture of a graphical “delivery tube” from the right display to the center; she then picks up the table’s customized product collection, drags it up to the center screen, and deposits it in the mouth of the delivery tube. The tube ingests the collection and retracts back to the right display; the collection will be delivered to the operator’s colleague, who is expecting to review her work and use it to construct an interactive visualization of a pet shop aisle.

The MMID of an alternative embodiment includes a housing having a rectangular form-factor. The pointer of this alternative embodiment is five inches long, one and one half inches wide, and one half inch deep, for example, but many other sizes and/or configurations are possible hereunder. The MMID includes optically tracked tags, described in detail below. The MMID does not include electronics as the processing software runs in a host system environment, but the embodiment is not so limited.

A user most naturally holds the pointer such that the long axis serves to point at objects (including virtual objects) in the user’s environment. The pointer can be rotated around the long axis to transition between two modal orientations (e.g., modes i and ii). Four modal transitions are possible, even though there are only two modes, because the system can distinguish between the direction of rotation during a transition: transition from mode i to mode ii/clockwise; transition from mode i to mode ii/counter-clockwise; transition from mode ii to mode i/clockwise; transition from mode ii to mode i/counter-clockwise. As with the MMID described above, these rotational transitions are tracked in input processing software, and can be subject to hysteretic locking.

The optical tags are mounted on the “front” portion (e.g., front half) of the pointer, in the area extending outwards from the user’s hand, for example, but are not so limited. On each of the two sides of the pointer, two tags are mounted. The forward-most tag on each side is fixed in position. The rear-most tag on each side is positioned a distance (e.g., five (5) centimeters) behind the forward tag and is aligned along and oriented according to the same axis. This rear tag is affixed to a spring-mounted sliding mechanism (the direction of translation aligned with the pointer’s long axis) such that the user’s thumb may push forward on the mechanism to decrease the distance between the two tags by approximately one centimeter.

The input processing software interprets the logical button state of the device to be in state (0) when the distance between the two tags is five centimeters. To effect a transition to state (1), the rear tag is moved a distance closer to the front tag (e.g., to within 4.2 centimeters of the front tag). The transition back to button state (1) is triggered only when the distance between the tags exceeds 4.8 centimeters. This is similar to the hysteresis applied to the device’s principal (rotational) mode transitions. Again, the size of the hysteresis band is configurable.

In the embodiment of an optically tracked MMID, an optical tracking tag is used where a number of dots are aligned on

a tag. These dots may be small spheres covered with retroreflectors, for example, allowing an IR tracking system (described below) to determine the location and orientation of a tagged object. In the case that this tagged object is an input MMID, it may be desired to provide a means for the tracking system to determine when a user has provided a non-geometric, state-change input, such as pressing a button.

The MMID of various alternative embodiments operates using infrared (IR) light-emitting diodes (LEDs) (IR LEDs) to provide tracking dots that are only visible to a camera at certain states based on the user input. The MMID of these alternative embodiments includes a battery and LED driving circuitry controlled by the input button. FIGS. 4a and 4b show input states of the MMID with IR LEDs, under an embodiment. The tag of this embodiment comprises numerous retro-reflective dots 402 (shown as a solid filled dot) and two IR LEDs 403 and 404. In FIG. 4a, the tag is shown in a state in which the button on the MMID is not pressed, and IR LED 403 is in the non-illuminated state, while IR LED 404 is in the illuminated state. In FIG. 4b, the user has pressed a button on the MMID and, in response, IR LED 403 is in the illuminated state while IR LED 404 is in the non-illuminated state. The optical processing system detects the difference in the two tags and from the state of the two tags determines the user’s intent.

FIGS. 5a and 5b show input states of the MMID with IR LEDs, under another alternative embodiment. In this embodiment, only one LED is switched. Thus, referring to FIG. 5a, LED 504 is in the non-illuminated state when the user has not pressed the button. In FIG. 5b, the user has pressed the button and LED 504 is thus illuminated.

Additional methods are also enabled using similar approaches. In one alternative embodiment, a complete tag is constructed using LEDs and the presence or absence of that tag provides input of the user. In another embodiment, two identical tags are created either overlaid (offset by, for example 0.5 cm) or adjacent. Illuminating one tag or the other, and determining the location of that tag with respect to another tag, allows the input state of the user to be determined.

The MMID of other alternative embodiments can combine the use of tag tracking with EMF tracking. These alternative embodiments combine aspects of the EMF tracking with the tag tracking using various types of tags, as described herein.

The MMID of another alternative embodiment includes a controller used in conjunction with two infrared light sources, one located in front of the user and one positioned behind the user. These two light sources each have three individual infrared emitters, and the emitter of each source is configured in a different pattern. The MMID of this embodiment makes use of inertial tracking, includes two modes, and includes multiple mechanical input buttons, as described below.

The MMID of this embodiment might be thought of as a modification of a Nintendo® Wii™ remote control device that supports two modal orientations, with the modes determined by the directional orientation of the controller relative to its environment. The Wii™ controller is a small device used to play video games on the Nintendo® Wii™ platform, and an associated infrared light source. The controller tracks its motion in space inertially, using a set of low-accuracy accelerometers. The accelerometers are not accurate enough to provide good position and orientation data over more than a few tenths of seconds, because of the errors that accumulate during numerical integration, so an optical tracking system (in conjunction with the light source component) is also used. The optical tracking system of the Wii™ controller therefore further comprises an internal, front-facing infrared camera

capable of locating four bright infrared light sources in a two-dimensional image plane. Therefore, the camera is embedded in the tracked device and the objects that are optically located are fixed-position environmental referents. By measuring the perceived size and position of known infrared light sources in the environment it is possible to determine the direction in which the controller is pointing and to triangulate the controllers distance from those sources. This infrared tracking technology may be viewed as an inversion of the tracking technology described herein, because the infrared tracking technology of the embodiment herein uses cameras placed in the environment to optically locate points arranged on devices, surfaces, gloves, and other objects.

In a typical use with the Nintendo® Wii™ console, the controller is always pointing towards a display screen. An infrared light source is placed above or below the display screen, providing the controller with a screen-relative orientation. In contrast, the controller of an embodiment is used in conjunction with two infrared light sources, one positioned in front of the user and one positioned behind the user. These two light sources each have three individual infrared emitters, and each source's emitters are configured in a different pattern.

The controller of an embodiment communicates by blue-tooth radio with input processing software or components running on a host computer. The input processing software identifies which emitter pattern is detected and therefore whether the controller is pointing forwards or backwards. Two modal orientations are derived from this forwards/backwards determination. In modal state (i) the controller is oriented forwards. In modal state (ii) the controller is oriented backwards. In each case, the user is logically pointing forwards. The user controls the mode by turning the controller around "back to front". This is in contrast to the embodiments described above, in which the mode control is a long-axis "rolling" of the device. The controller of an embodiment can include an embedded speaker, providing sound output, several lights, and a vibration (or "rumble") output.

Numerous modifications of the embodiments described herein are possible under this description. The controller of an embodiment may, for example, have two cameras, one on each end of the device, thereby obviating the need for two light sources. The light sources may be differentiated by timing, rather than spatial, patterns.

Spatial Operating Environment (SOE)

Embodiments of a spatial-continuum input system are described herein in the context of a Spatial Operating Environment (SOE). As an example, FIG. 6 is a block diagram of a Spatial Operating Environment (SOE), under an embodiment. A user locates his hands **101** and **102** in the viewing area **150** of an array of cameras **104A-104D**. The cameras detect location, orientation, and movement of the fingers and hands **101** and **102**, as spatial tracking data, and generate output signals to pre-processor **105**. Pre-processor **105** translates the camera output into a gesture signal that is provided to the computer processing unit **107** of the system. The computer **107** uses the input information to generate a command to control one or more on screen cursors and provides video output to display **103**.

Although the system is shown with a single user's hands as input, the SOE **100** may be implemented using multiple users. In addition, instead of or in addition to hands, the system may track any part or parts of a user's body, including head, feet, legs, arms, elbows, knees, and the like.

In the embodiment shown, four cameras or sensors are used to detect the location, orientation, and movement of the user's hands **101** and **102** in the viewing area **150**. It should be

understood that the SOE **100** may include more (e.g., six cameras, eight cameras, etc.) or fewer (e.g., two cameras) cameras or sensors without departing from the scope or spirit of the SOE. In addition, although the cameras or sensors are disposed symmetrically in the example embodiment, there is no requirement of such symmetry in the SOE **100**. Any number or positioning of cameras or sensors that permits the location, orientation, and movement of the user's hands may be used in the SOE **100**.

In one embodiment, the cameras used are motion capture cameras capable of capturing grey-scale images. In one embodiment, the cameras used are those manufactured by Vicon, such as the Vicon MX40 camera. This camera includes on-camera processing and is capable of image capture at 1000 frames per second. A motion capture camera is capable of detecting and locating markers.

In the embodiment described, the cameras are sensors used for optical detection. In other embodiments, the cameras or other detectors may be used for electromagnetic, magneto-static, RFID, or any other suitable type of detection.

Pre-processor **105** generates three dimensional space point reconstruction and skeletal point labeling. The gesture translator **106** converts the 3D spatial information and marker motion information into a command language that can be interpreted by a computer processor to update the location, shape, and action of a cursor on a display. In an alternate embodiment of the SOE **100**, the pre-processor **105** and gesture translator **106** are integrated or combined into a single device.

Computer **107** may be any general purpose computer such as manufactured by Apple, Dell, or any other suitable manufacturer. The computer **107** runs applications and provides display output. Cursor information that would otherwise come from a mouse or other prior art input device now comes from the gesture system.

Marker Tags

The SOE or an embodiment contemplates the use of marker tags on one or more fingers of the user so that the system can locate the hands of the user, identify whether it is viewing a left or right hand, and which fingers are visible. This permits the system to detect the location, orientation, and movement of the user's hands. This information allows a number of gestures to be recognized by the system and used as commands by the user.

The marker tags in one embodiment are physical tags comprising a substrate (appropriate in the present embodiment for affixing to various locations on a human hand) and discrete markers arranged on the substrate's surface in unique identifying patterns.

The markers and the associated external sensing system may operate in any domain (optical, electromagnetic, magnetostatic, etc.) that allows the accurate, precise, and rapid and continuous acquisition of their three-space position. The markers themselves may operate either actively (e.g. by emitting structured electromagnetic pulses) or passively (e.g. by being optically retroreflective, as in the present embodiment).

At each frame of acquisition, the detection system receives the aggregate 'cloud' of recovered three-space locations comprising all markers from tags presently in the instrumented workspace volume (within the visible range of the cameras or other detectors). The markers on each tag are of sufficient multiplicity and are arranged in unique patterns such that the detection system can perform the following tasks: (1) segmentation, in which each recovered marker position is assigned to one and only one subcollection of points that form a single tag; (2) labelling, in which each segmented subcollection of points is identified as a particular tag; (3) location,

in which the three-space position of the identified tag is recovered; and (4) orientation, in which the three-space orientation of the identified tag is recovered. Tasks (1) and (2) are made possible through the specific nature of the marker-patterns, as described below and as illustrated in one embodiment in FIG. 7.

The markers on the tags in one embodiment are affixed at a subset of regular grid locations. This underlying grid may, as in the present embodiment, be of the traditional Cartesian sort; or may instead be some other regular plane tessellation (a triangular/hexagonal tiling arrangement, for example). The scale and spacing of the grid is established with respect to the known spatial resolution of the marker-sensing system, so that adjacent grid locations are not likely to be confused. Selection of marker patterns for all tags should satisfy the following constraint: no tag's pattern shall coincide with that of any other tag's pattern through any combination of rotation, translation, or mirroring. The multiplicity and arrangement of markers may further be chosen so that loss (or occlusion) of some specified number of component markers is tolerated: After any arbitrary transformation, it should still be unlikely to confuse the compromised module with any other.

Referring now to FIG. 7, a number of tags **201A-201E** (left hand) and **202A-202E** (right hand) are shown. Each tag is rectangular and consists in this embodiment of a 5×7 grid array. The rectangular shape is chosen as an aid in determining orientation of the tag and to reduce the likelihood of mirror duplicates. In the embodiment shown, there are tags for each finger on each hand. In some embodiments, it may be adequate to use one, two, three, or four tags per hand. Each tag has a border of a different grey-scale or color shade. Within this border is a 3×5 grid array. Markers (represented by the black dots of FIG. 7) are disposed at certain points in the grid array to provide information.

Qualifying information may be encoded in the tags' marker patterns through segmentation of each pattern into 'common' and 'unique' subpatterns. For example, the present embodiment specifies two possible 'border patterns', distributions of markers about a rectangular boundary. A 'family' of tags is thus established—the tags intended for the left hand might thus all use the same border pattern as shown in tags **201A-201E** while those attached to the right hand's fingers could be assigned a different pattern as shown in tags **202A-202E**. This subpattern is chosen so that in all orientations of the tags, the left pattern can be distinguished from the right pattern. In the example illustrated, the left hand pattern includes a marker in each corner and on marker in a second from corner grid location. The right hand pattern has markers in only two corners and two markers in non corner grid locations. An inspection of the pattern reveals that as long as any three of the four markers are visible, the left hand pattern can be positively distinguished from the left hand pattern. In one embodiment, the color or shade of the border can also be used as an indicator of handedness.

Each tag must of course still employ a unique interior pattern, the markers distributed within its family's common border. In the embodiment shown, it has been found that two markers in the interior grid array are sufficient to uniquely identify each of the ten fingers with no duplication due to rotation or orientation of the fingers. Even if one of the markers is occluded, the combination of the pattern and the handedness of the tag yields a unique identifier.

In the present embodiment, the grid locations are visually present on the rigid substrate as an aid to the (manual) task of affixing each retroreflective marker at its intended location. These grids and the intended marker locations are literally printed via color inkjet printer onto the substrate, which here

is a sheet of (initially) flexible 'shrink-film'. Each module is cut from the sheet and then oven-baked, during which thermal treatment each module undergoes a precise and repeatable shrinkage. For a brief interval following this procedure, the cooling tag may be shaped slightly—to follow the longitudinal curve of a finger, for example; thereafter, the substrate is suitably rigid, and markers may be affixed at the indicated grid points.

In one embodiment, the markers themselves are three dimensional, such as small reflective spheres affixed to the substrate via adhesive or some other appropriate means. The three-dimensionality of the markers can be an aid in detection and location over two dimensional markers. However either can be used without departing from the spirit and scope of the SOE described herein.

At present, tags are affixed via Velcro or other appropriate means to a glove worn by the operator or are alternately affixed directly to the operator's fingers using a mild double-stick tape. In a third embodiment, it is possible to dispense altogether with the rigid substrate and affix—or 'paint'—individual markers directly onto the operator's fingers and hands.

Gesture Vocabulary

The SOE of an embodiment contemplates a gesture vocabulary consisting of hand poses, orientation, hand combinations, and orientation blends. A notation language is also implemented for designing and communicating poses and gestures in the gesture vocabulary of the SOE. The gesture vocabulary is a system for representing instantaneous 'pose states' of kinematic linkages in compact textual form. The linkages in question may be biological (a human hand, for example; or an entire human body; or a grasshopper leg; or the articulated spine of a lemur) or may instead be nonbiological (e.g. a robotic arm). In any case, the linkage may be simple (the spine) or branching (the hand). The gesture vocabulary system of the SOE establishes for any specific linkage a constant length string; the aggregate of the specific ASCII characters occupying the string's 'character locations' is then a unique description of the instantaneous state, or 'pose', of the linkage.

Hand Poses

FIG. 8 illustrates hand poses in an embodiment of a gesture vocabulary of the SOE, under an embodiment. The SOE supposes that each of the five fingers on a hand is used. These fingers are codes as p-pinkie, r-ring finger, m-middle finger, i-index finger, and t-thumb. A number of poses for the fingers and thumbs are defined and illustrated in FIG. 8. A gesture vocabulary string establishes a single character position for each expressible degree of freedom in the linkage (in this case, a finger). Further, each such degree of freedom is understood to be discretized (or 'quantized'), so that its full range of motion can be expressed through assignment of one of a finite number of standard ASCII characters at that string position. These degrees of freedom are expressed with respect to a body-specific origin and coordinate system (the back of the hand, the center of the grasshopper's body; the base of the robotic arm; etc.). A small number of additional gesture vocabulary character positions are therefore used to express the position and orientation of the linkage 'as a whole' in the more global coordinate system.

Still referring to FIG. 8, a number of poses are defined and identified using ASCII characters. Some of the poses are divided between thumb and non-thumb. The SOE in this embodiment uses a coding such that the ASCII character itself is suggestive of the pose. However, any character may be used to represent a pose, whether suggestive or not. In addition, there is no requirement in the embodiments to use ASCII

characters for the notation strings. Any suitable symbol, numeral, or other representation maybe used without departing from the scope and spirit of the embodiments. For example, the notation may use two bits per finger if desired or some other number of bits as desired.

A curled finger is represented by the character “^” while a curled thumb by “>”. A straight finger or thumb pointing up is indicated by “1” and at an angle by “\” or “/”. “-” represents a thumb pointing straight sideways and “x” represents a thumb pointing into the plane.

Using these individual finger and thumb descriptions, a robust number of hand poses can be defined and written using the scheme of the embodiments. Each pose is represented by five characters with the order being p-r-m-i-t as described above. FIG. 8 illustrates a number of poses and a few are described here by way of illustration and example. The hand held flat and parallel to the ground is represented by “11111”. A fist is represented by “^^^>”. An “OK” sign is represented by “111^>”

The character strings provide the opportunity for straight-forward ‘human readability’ when using suggestive characters. The set of possible characters that describe each degree of freedom may generally be chosen with an eye to quick recognition and evident analogy. For example, a vertical bar (‘|’) would likely mean that a linkage element is ‘straight’, an ell (‘L’) might mean a ninety-degree bend, and a circumflex (‘^’) could indicate a sharp bend. As noted above, any characters or coding may be used as desired.

Any system employing gesture vocabulary strings such as described herein enjoys the benefit of the high computational efficiency of string comparison—identification of or search for any specified pose literally becomes a ‘string compare’ (e.g. UNIX’s strep() function) between the desired pose string and the instantaneous actual string. Furthermore, the use of ‘wildcard characters’ provides the programmer or system designer with additional familiar efficiency and efficacy: degrees of freedom whose instantaneous state is irrelevant for a match may be specified as an interrogation point (?); additional wildcard meanings may be assigned.

Orientation

In addition to the pose of the fingers and thumb, the orientation of the hand can represent information. Characters describing global-space orientations can also be chosen transparently: the characters ‘<’, ‘>’, ‘^’, and ‘v’ may be used to indicate, when encountered in an orientation character position, the ideas of left, right, up, and down. FIG. 9 illustrates hand orientation descriptors and examples of coding that combines pose and orientation. In an embodiment, two character positions specify first the direction of the palm and then the direction of the fingers (if they were straight, irrespective of the fingers’ actual bends). The possible characters for these two positions express a ‘body-centric’ notion of orientation: ‘-’, ‘+’, ‘x’, ‘*’, ‘^’, and ‘v’ describe medial, lateral, anterior (forward, away from body), posterior (backward, away from body), cranial (upward), and caudal (downward).

In the notation scheme of an embodiment, the five finger pose indicating characters are followed by a colon and then two orientation characters to define a complete command pose. In one embodiment, a start position is referred to as an “xyz” pose where the thumb is pointing straight up, the index finger is pointing forward and the middle finger is perpendicular to the index finger, pointing to the left when the pose is made with the right hand. This is represented by the string “^^x1-:-x”.

‘XYZ-hand’ is a technique for exploiting the geometry of the human hand to allow full six-degree-of-freedom navigation of visually presented three-dimensional structure.

Although the technique depends only on the bulk translation and rotation of the operator’s hand—so that its fingers may in principal be held in any pose desired—the present embodiment prefers a static configuration in which the index finger points away from the body; the thumb points toward the ceiling; and the middle finger points left-right. The three fingers thus describe (roughly, but with clearly evident intent) the three mutually orthogonal axes of a three-space coordinate system: thus ‘XYZ-hand’.

XYZ-hand navigation then proceeds with the hand, fingers in a pose as described above, held before the operator’s body at a predetermined ‘neutral location’. Access to the three translational and three rotational degrees of freedom of a three-space object (or camera) is effected in the following natural way: left-right movement of the hand (with respect to the body’s natural coordinate system) results in movement along the computational context’s x-axis; up-down movement of the hand results in movement along the controlled context’s y-axis; and forward-back hand movement (toward/away from the operator’s body) results in z-axis motion within the context. Similarly, rotation of the operator’s hand about the index finger leads to a ‘roll’ change of the computational context’s orientation; ‘pitch’ and ‘yaw’ changes are effected analogously, through rotation of the operator’s hand about the middle finger and thumb, respectively.

Note that while ‘computational context’ is used here to refer to the entity being controlled by the XYZ-hand method—and seems to suggest either a synthetic three-space object or camera—it should be understood that the technique is equally useful for controlling the various degrees of freedom of real-world objects: the pan/tilt/roll controls of a video or motion picture camera equipped with appropriate rotational actuators, for example. Further, the physical degrees of freedom afforded by the XYZ-hand posture may be somewhat less literally mapped even in a virtual domain: In the present embodiment, the XYZ-hand is also used to provide navigational access to large panoramic display images, so that left-right and up-down motions of the operator’s hand lead to the expected left-right or up-down ‘panning’ about the image, but forward-back motion of the operator’s hand maps to ‘zooming’ control.

In every case, coupling between the motion of the hand and the induced computational translation/rotation may be either direct (i.e. a positional or rotational offset of the operator’s hand maps one-to-one, via some linear or nonlinear function, to a positional or rotational offset of the object or camera in the computational context) or indirect (i.e. positional or rotational offset of the operator’s hand maps one-to-one, via some linear or nonlinear function, to a first or higher-degree derivative of position/orientation in the computational context; ongoing integration then effects a non-static change in the computational context’s actual zero-order position/orientation). This latter means of control is analogous to use of an automobile’s ‘gas pedal’, in which a constant offset of the pedal leads, more or less, to a constant vehicle speed.

The ‘neutral location’ that serves as the real-world XYZ-hand’s local six-degree-of-freedom coordinate origin may be established (1) as an absolute position and orientation in space (relative, say, to the enclosing room); (2) as a fixed position and orientation relative to the operator herself (e.g. eight inches in front of the body, ten inches below the chin, and laterally in line with the shoulder plane), irrespective of the overall position and ‘heading’ of the operator; or (3) interactively, through deliberate secondary action of the operator (using, for example, a gestural command enacted by the operator’s ‘other’ hand, said command indicating that the

XYZ-hand's present position and orientation should henceforth be used as the translational and rotational origin).

It is further convenient to provide a 'detent' region (or 'dead zone') about the XYZ-hand's neutral location, such that movements within this volume do not map to movements in the controlled context.

Other poses may included:

[||||:vx] is a flat hand (thumb parallel to fingers) with palm facing down and fingers forward.

[||||:x^] is a flat hand with palm facing forward and fingers toward ceiling.

[||||:-x] is a flat hand with palm facing toward the center of the body (right if left hand, left if right hand) and fingers forward.

[^^^:-x] is a single-hand thumbs-up (with thumb pointing toward ceiling)

[^^|-x] is a mime gun pointing forward.

Two Hand Combination

The SOE of an embodiment contemplates single hand commands and poses, as well as two-handed commands and poses. FIG. 10 illustrates examples of two hand combinations and associated notation in an embodiment of the SOE. Reviewing the notation of the first example, "full stop" reveals that it comprises two closed fists. The "snapshot" example has the thumb and index finger of each hand extended, thumbs pointing toward each other, defining a goal post shaped frame. The "rudder and throttle start position" is fingers and thumbs pointing up palms facing the screen.

Orientation Blends

FIG. 11 illustrates an example of an orientation blend in an embodiment of the SOE. In the example shown the blend is represented by enclosing pairs of orientation notations in parentheses after the finger pose string. For example, the first command shows finger positions of all pointing straight. The first pair of orientation commands would result in the palms being flat toward the display and the second pair has the hands rotating to a 45 degree pitch toward the screen. Although pairs of blends are shown in this example, any number of blends is contemplated in the SOE.

Example Commands

FIGS. 13/1 and 13/2 show a number of possible commands that may be used with the SOE. Although some of the discussion here has been about controlling a cursor on a display, the SOE is not limited to that activity. In fact, the SOE has great application in manipulating any and all data and portions of data on a screen, as well as the state of the display. For example, the commands may be used to take the place of video controls during play back of video media. The commands may be used to pause, fast forward, rewind, and the like. In addition, commands may be implemented to zoom in or zoom out of an image, to change the orientation of an image, to pan in any direction, and the like. The SOE may also be used in lieu of menu commands such as open, close, save, and the like. In other words, any commands or activity that can be imagined can be implemented with hand gestures.

Operation

FIG. 12 is a flow diagram illustrating the operation of the SOE in one embodiment. At 701 the detection system detects the markers and tags. At 702 it is determined if the tags and markers are detected. If not, the system returns to 701. If the tags and markers are detected at 702, the system proceeds to 703. At 703 the system identifies the hand, fingers and pose from the detected tags and markers. At 704 the system identifies the orientation of the pose. At 705 the system identifies the three dimensional spatial location of the hand or hands that are detected. (Please note that any or all of 703, 704, and 705 may be combined).

At 706 the information is translated to the gesture notation described above. At 707 it is determined if the pose is valid. This may be accomplished via a simple string comparison using the generated notation string. If the pose is not valid, the system returns to 701. If the pose is valid, the system sends the notation and position information to the computer at 708. At 709 the computer determines the appropriate action to take in response to the gesture and updates the display accordingly at 710.

In one embodiment of the SOE, 701-705 are accomplished by the on-camera processor. In other embodiments, the processing can be accomplished by the system computer if desired.

Parsing and Translation

The system is able to "parse" and "translate" a stream of low-level gestures recovered by an underlying system, and turn those parsed and translated gestures into a stream of command or event data that can be used to control a broad range of computer applications and systems. These techniques and algorithms may be embodied in a system consisting of computer code that provides both an engine implementing these techniques and a platform for building computer applications that make use of the engine's capabilities.

One embodiment is focused on enabling rich gestural use of human hands in computer interfaces, but is also able to recognize gestures made by other body parts (including, but not limited to arms, torso, legs and the head), as well as non-hand physical tools of various kinds, both static and articulating, including but not limited to calipers, compasses, flexible curve approximators, and pointing devices of various shapes. The markers and tags may be applied to items and tools that may be carried and used by the operator as desired.

The system described here incorporates a number of innovations that make it possible to build gestural systems that are rich in the range of gestures that can be recognized and acted upon, while at the same time providing for easy integration into applications.

The gestural parsing and translation system in one embodiment comprises:

1) a compact and efficient way to specify (encode for use in computer programs) gestures at several different levels of aggregation:

- a. a single hand's "pose" (the configuration and orientation of the parts of the hand relative to one another) a single hand's orientation and position in three-dimensional space.
- b. two-handed combinations, for either hand taking into account pose, position or both.
- c. multi-person combinations; the system can track more than two hands, and so more than one person can cooperatively (or competitively, in the case of game applications) control the target system.
- d. sequential gestures in which poses are combined in a series; we call these "animating" gestures.
- e. "grapheme" gestures, in which the operator traces shapes in space.

2) a programmatic technique for registering specific gestures from each category above that are relevant to a given application context.

3) algorithms for parsing the gesture stream so that registered gestures can be identified and events encapsulating those gestures can be delivered to relevant application contexts.

The specification system (1), with constituent elements (1a) to (1f), provides the basis for making use of the gestural parsing and translating capabilities of the system described here.

A single-hand “pose” is represented as a string of

i) relative orientations between the fingers and the back of the hand,

ii) quantized into a small number of discrete states.

Using relative joint orientations allows the system described here to avoid problems associated with differing hand sizes and geometries. No “operator calibration” is required with this system. In addition, specifying poses as a string or collection of relative orientations allows more complex gesture specifications to be easily created by combining pose representations with further filters and specifications.

Using a small number of discrete states for pose specification makes it possible to specify poses compactly as well as to ensure accurate pose recognition using a variety of underlying tracking technologies (for example, passive optical tracking using cameras, active optical tracking using lighted dots and cameras, electromagnetic field tracking, etc).

Gestures in every category (1a) to (1f) may be partially (or minimally) specified, so that non-critical data is ignored. For example, a gesture in which the position of two fingers is definitive, and other finger positions are unimportant, may be represented by a single specification in which the operative positions of the two relevant fingers is given and, within the same string, “wild cards” or generic “ignore these” indicators are listed for the other fingers.

All of the innovations described here for gesture recognition, including but not limited to the multi-layered specification technique, use of relative orientations, quantization of data, and allowance for partial or minimal specification at every level, generalize beyond specification of hand gestures to specification of gestures using other body parts and “manufactured” tools and objects.

The programmatic techniques for “registering gestures” (2), consist of a defined set of Application Programming Interface calls that allow a programmer to define which gestures the engine should make available to other parts of the running system.

These API routines may be used at application set-up time, creating a static interface definition that is used throughout the lifetime of the running application. They may also be used during the course of the run, allowing the interface characteristics to change on the fly. This real-time alteration of the interface makes it possible to,

i) build complex contextual and conditional control states,

ii) to dynamically add hysteresis to the control environment, and

iii) to create applications in which the user is able to alter or extend the interface vocabulary of the running system itself.

Algorithms for parsing the gesture stream (3) compare gestures specified as in (1) and registered as in (2) against incoming low-level gesture data. When a match for a registered gesture is recognized, event data representing the matched gesture is delivered up the stack to running applications.

Efficient real-time matching is desired in the design of this system, and specified gestures are treated as a tree of possibilities that are processed as quickly as possible.

In addition, the primitive comparison operators used internally to recognize specified gestures are also exposed for the applications programmer to use, so that further comparison (flexible state inspection in complex or compound gestures, for example) can happen even from within application contexts.

Recognition “locking” semantics are an innovation of the system described here. These semantics are implied by the registration API (2) (and, to a lesser extent, embedded within the specification vocabulary (1)). Registration API calls include,

i) “entry” state notifiers and “continuation” state notifiers, and

ii) gesture priority specifiers.

If a gesture has been recognized, its “continuation” conditions take precedence over all “entry” conditions for gestures of the same or lower priorities. This distinction between entry and continuation states adds significantly to perceived system usability.

The system described here includes algorithms for robust operation in the face of real-world data error and uncertainty. Data from low-level tracking systems may be incomplete (for a variety of reasons, including occlusion of markers in optical tracking, network drop-out or processing lag, etc).

Missing data is marked by the parsing system, and interpolated into either “last known” or “most likely” states, depending on the amount and context of the missing data.

If data about a particular gesture component (for example, the orientation of a particular joint) is missing, but the “last known” state of that particular component can be analyzed as physically possible, the system uses this last known state in its real-time matching.

Conversely, if the last known state is analyzed as physically impossible, the system falls back to a “best guess range” for the component, and uses this synthetic data in its real-time matching.

The specification and parsing systems described here have been carefully designed to support “handedness agnosticism,” so that for multi-hand gestures either hand is permitted to satisfy pose requirements.

35 Navigating Data Space

The SOE of an embodiment enables ‘pushback’, a linear spatial motion of a human operator’s hand, or performance of analogously dimensional activity, to control linear verging or trucking motion through a graphical or other data-representational space. The SOE, and the computational and cognitive association established by it, provides a fundamental, structured way to navigate levels of scale, to traverse a principally linear ‘depth dimension’, or—most generally—to access quantized or ‘detented’ parameter spaces. The SOE also provides an effective means by which an operator may volitionally acquire additional context: a rapid technique for understanding vicinities and neighborhoods, whether spatial, conceptual, or computational.

In certain embodiments, the pushback technique may employ traditional input devices (e.g. mouse, trackball, integrated sliders or knobs) or may depend on tagged or tracked objects external to the operator’s own person (e.g. instrumented kinematic linkages, magnetostatically tracked ‘input bricks’). In other alternative embodiments, a pushback implementation may suffice as the whole of a control system.

The SOE of an embodiment is a component of and integrated into a larger spatial interaction system that supplants customary mouse-based graphical user interface (‘WIMP’ UI) methods for control of a computer, comprising instead (a) physical sensors that can track one or more types of object (e.g., human hands, objects on human hands, inanimate objects, etc.); (b) an analysis component for analyzing the evolving position, orientation, and pose of the sensed hands into a sequence of gestural events; (c) a descriptive scheme for representing such spatial and gestural events; (d) a framework for distributing such events to and within control programs; (e) methods for synchronizing the human intent (the

commands) encoded by the stream of gestural events with graphical, aural, and other display-modal depictions of both the event stream itself and of the application-specific consequences of event interpretation, all of which are described in detail below. In such an embodiment, the pushback system is integrated with additional spatial and gestural input-and-interface techniques.

Generally, the navigation of a data space comprises detecting a gesture of a body from gesture data received via a detector. The gesture data is absolute three-space location data of an instantaneous state of the body at a point in time and physical space. The detecting comprises identifying the gesture using the gesture data. The navigating comprises translating the gesture to a gesture signal, and navigating through the data space in response to the gesture signal. The data space is a data-representational space comprising a dataset represented in the physical space.

When an embodiment's overall round-trip latency (hand motion to sensors to pose analysis to pushback interpretation system to computer graphics rendering to display device back to operator's visual system) is kept low (e.g., an embodiment exhibits latency of approximately fifteen milliseconds) and when other parameters of the system are properly tuned, the perceptual consequence of pushback interaction is a distinct sense of physical causality: the SOE literalizes the physically resonant metaphor of pushing against a spring-loaded structure. The perceived causality is a highly effective feedback; along with other more abstract graphical feedback modalities provided by the pushback system, and with a deliberate suppression of certain degrees of freedom in the interpretation of operator movement, such feedback in turn permits stable, reliable, and repeatable use of both gross and fine human motor activity as a control mechanism.

In evaluating the context of the SOE, many datasets are inherently spatial: they represent phenomena, events, measurements, observations, or structure within a literal physical space. For other datasets that are more abstract or that encode literal yet non-spatial information, it is often desirable to prepare a representation (visual, aural, or involving other display modalities) some fundamental aspect of which is controlled by a single, scalar-valued parameter; associating that parameter with a spatial dimension is then frequently also beneficial. It is manipulation of this single scalar parameter, as is detailed below, which benefits from manipulation by means of the pushback mechanism.

Representations may further privilege a small plurality of discrete values of their parameter—indeed, sometimes only one—at which the dataset is optimally regarded. In such cases it is useful to speak of a 'detented parameter' or, if the parameter has been explicitly mapped onto one dimension of a representational space, of 'detented space'. Use of the term 'detented' herein is intended to evoke not only the preferential quantization of the parameter but also the visuo-haptic sensation of ratchets, magnetic alignment mechanisms, jog-shuttle wheels, and the wealth of other worldly devices that are possessed of deliberate mechanical detents.

Self-evident yet crucially important examples of such parameters include but are not limited to (1) the distance of a synthetic camera, in a computer graphics environment, from a renderable representation of a dataset; (2) the density at which data is sampled from the original dataset and converted into renderable form; (3) the temporal index at which samples are retrieved from a time-varying dataset and converted to a renderable representation. These are universal approaches; countless domain-specific parameterizations also exist.

The pushback of the SOE generally aligns the dataset's parameter-control axis with a locally relevant 'depth dimen-

sion' in physical space, and allows structured real-world motion along the depth dimension to effect a data-space translation along the control axis. The result is a highly efficient means for navigating a parameter space. Following are detailed descriptions of representative embodiments of the pushback as implemented in the SOE.

In a pushback example, an operator stands at a comfortable distance before a large wall display on which appears a single 'data frame' comprising text and imagery, which graphical data elements may be static or dynamic. The data frame, for example, can include an image, but is not so limited. The data frame, itself a two-dimensional construct, is nonetheless resident in a three-dimensional computer graphics rendering environment whose underlying coordinate system has been arranged to coincide with real-world coordinates convenient for describing the room and its contents, including the display and the operator.

The operator's hands are tracked by sensors that resolve the position and orientation of her fingers, and possibly of the overall hand masses, to high precision and at a high temporal rate; the system analyzes the resulting spatial data in order to characterize the 'pose' of each hand—i.e. the geometric disposition of the fingers relative to each other and to the hand mass. While this example embodiment tracks an object that is a human hand(s), numerous other objects could be tracked as input devices in alternative embodiments. One example is a one-sided pushback scenario in which the body is an operator's hand in the open position, palm facing in a forward direction (along the z-axis) (e.g., toward a display screen in front of the operator). For the purposes of this description, the wall display is taken to occupy the x and y dimensions; z describes the dimension between the operator and the display. The gestural interaction space associated with this pushback embodiment comprises two spaces abutted at a plane of constant z; the detented interval space farther from the display (i.e. closer to the operator) is termed the 'dead zone', while the closer half-space is the 'active zone'. The dead zone extends indefinitely in the backward direction (toward the operator and away from the display) but only a finite distance forward, ending at the dead zone threshold. The active zone extends from the dead zone threshold forward to the display. The data frame(s) rendered on the display are interactively controlled or "pushed back" by movements of the body in the active zone.

The data frame is constructed at a size and aspect ratio precisely matching those of the display, and is positioned and oriented so that its center and normal vector coincide with those physical attributes of the display, although the embodiment is not so limited. The virtual camera used to render the scene is located directly forward from the display and at roughly the distance of the operator. In this context, the rendered frame thus precisely fills the display.

Arranged logically to the left and right of the visible frame are a number of additional coplanar data frames, uniformly spaced and with a modest gap separating each from its immediate neighbors. Because they lie outside the physical/virtual rendering bounds of the computer graphics rendering geometry, these laterally displaced adjacent data frames are not initially visible. As will be seen, the data space—given its geometric structure—is possessed of a single natural detent in the z-direction and a plurality of x-detents.

The operator raises her left hand, held in a loose fist pose, to her shoulder. She then extends the fingers so that they point upward and the thumb so that it points to the right; her palm faces the screen (in the gestural description language described in detail below, this pose transition would be expressed as [^^^>:x^into |||-x^]). The system, detecting the

new pose, triggers pushback interaction and immediately records the absolute three-space hand position at which the pose was first entered: this position is used as the ‘origin’ from which subsequent hand motions will be reported as relative offsets.

Immediately, two concentric, partially transparent glyphs are superimposed on the center of the frame (and thus at the display’s center). For example, the glyphs can indicate body pushback gestures in the dead zone up to a point of the dead zone threshold. That the second glyph is smaller than the first glyph is an indication that the operator’s hand resides in the dead zone, through which the pushback operation is not ‘yet’ engaged. As the operator moves her hand forward (toward the dead zone threshold and the display), the second glyph incrementally grows. The second glyph is equivalent in size to the first glyph at the point at which the operator’s hand is at the dead zone threshold. The glyphs of this example describe the evolution of the glyph’s concentric elements as the operator’s hand travels forward from its starting position toward the dead zone threshold separating the dead zone from the active zone. The inner “toothy” part of the glyph, for example, grows as the hand nears the threshold, and is arranged so that the radius of the inner glyph and (static) outer glyph precisely match as the hand reaches the threshold position.

The second glyph shrinks in size inside the first glyph as the operator moves her hand away from the dead zone threshold and away from the display, remaining however always concentric with the first glyph and centered on the display. Crucially, only the z-component of the operator’s hand motion is mapped into the glyph’s scaling; incidental x- and y-components of the hand motion make no contribution.

When the operator’s hand traverses the forward threshold of the dead zone, crossing into the active zone, the pushback mechanism is engaged. The relative z-position of the hand (measured from the threshold) is subjected to a scaling function and the resulting value is used to effect a z-axis displacement of the data frame and its lateral neighbors, so that the rendered image of the frame is seen to recede from the display; the neighboring data frames also then become visible, ‘filling in’ from the edges of the display space—the constant angular subtent of the synthetic camera geometrically ‘captures’ more of the plane in which the frames lie as that plane moves away from the camera. The z-displacement is continuously updated, so that the operator, pushing her hand toward the display and pulling it back toward herself, perceives the lateral collection of frames receding and verging in direct response to her movements

As an example of a first relative z-axis displacement of the data frame resulting from corresponding pushback, the rendered image of the data frame is seen to recede from the display and the neighboring data frames become visible, ‘filling in’ from the edges of the display space. The neighboring data frames, which include a number of additional coplanar data frames, are arranged logically to the left and right of the visible frame, uniformly spaced and with a modest gap separating each from its immediate neighbors. As an example of a second relative z-axis displacement of the data frame resulting from corresponding pushback, and considering the first relative z-axis displacement, and assuming further pushing of the operator’s hand (pushing further along the z-axis toward the display and away from the operator) from that pushing resulting in the first relative z-axis displacement, the rendered image of the frame is seen to further recede from the display so that additional neighboring data frames become visible, further ‘filling in’ from the edges of the display space.

The paired concentric glyphs, meanwhile, now exhibit a modified feedback: with the operator’s hand in the active

zone, the second glyph switches from scaling-based reaction to a rotational reaction in which the hand’s physical z-axis offset from the threshold is mapped into a positive (in-plane) angular offset. In an example of the glyphs indicating body pushback gestures in the dead zone beyond the point of the dead zone threshold (along the z-axis toward the display and away from the operator), the glyphs depict the evolution of the glyph once the operator’s hand has crossed the dead zone threshold—i.e. when the pushback mechanism has been actively engaged. The operator’s hand movements toward and away from the display are thus visually indicated by clockwise and anticlockwise rotation of the second glyph (with the first glyph, as before, providing a static reference state), such that the “toothy” element of the glyph rotates as a linear function of the hand’s offset from the threshold, turning linear motion into a rotational representation.

Therefore, in this example, an additional first increment of hand movement along the z-axis toward the display is visually indicated by an incremental clockwise rotation of the second glyph (with the first glyph, as before, providing a static reference state), such that the “toothy” element of the glyph rotates a first amount corresponding to a linear function of the hand’s offset from the threshold. An additional second increment of hand movement along the z-axis toward the display is visually indicated by an incremental clockwise rotation of the second glyph (with the first glyph, as before, providing a static reference state), such that the “toothy” element of the glyph rotates a second amount corresponding to a linear function of the hand’s offset from the threshold. Further, a third increment of hand movement along the z-axis toward the display is visually indicated by an incremental clockwise rotation of the second glyph (with the first glyph, as before, providing a static reference state), such that the “toothy” element of the glyph rotates a third amount corresponding to a linear function of the hand’s offset from the threshold.

In this sample application, a secondary dimensional sensitivity is engaged when the operator’s hand is in the active zone: lateral (x-axis) motion of the hand is mapped, again through a possible scaling function, to x-displacement of the horizontal frame sequence. If the scaling function is positive, the effect is one of positional ‘following’ of the operator’s hand, and she perceives that she is sliding the frames left and right. As an example of a lateral x-axis displacement of the data frame resulting from lateral motion of the body, the data frames slide from left to right such that particular data frames disappear or partially disappear from view via the left edge of the display space while additional data frames fill in from the right edge of the display space.

Finally, when the operator causes her hand to exit the palm-forward pose (by, e.g., closing the hand into a fist), the pushback interaction is terminated and the collection of frames is rapidly returned to its original z-detent (i.e. coplanar with the display). Simultaneously, the frame collection is laterally adjusted to achieve x-coincidence of a single frame with the display; which frame ends thus ‘display-centered’ is whichever was closest to the concentric glyphs’ center at the instant of pushback termination: the nearest x-detent. The glyph structure is here seen serving a second function, as a selection reticle, but the embodiment is not so limited. The z- and x-positions of the frame collection are typically allowed to progress to their final display-coincident values over a short time interval in order to provide a visual sense of ‘spring-loaded return’.

The pushback system as deployed in this example provides efficient control modalities for (1) acquiring cognitively valuable ‘neighborhood context’ by variably displacing an aggregate dataset along the direct visual sightline—the depth

dimension—thereby bringing more of the dataset into view (in exchange for diminishing the angular subtent of any given part of the dataset); (2) acquiring neighborhood context by variably displacing the laterally-arrayed dataset along its natural horizontal dimension, maintaining the angular subtent of any given section of data but trading the visibility of old data for that of new data, in the familiar sense of ‘scrolling’; (3) selecting discretized elements of the dataset through rapid and dimensionally-constrained navigation.

In another example of the pushback of an embodiment, an operator stands immediately next to a waist-level display device whose active surface lies in a horizontal plane parallel to the floor. The coordinate system is here established in a way consistent with that of the previous example: the display surface lies in the x-z plane, so that the y-axis, representing the normal to the surface, is aligned in opposition to the physical gravity vector.

In an example physical scenario in which the body is held horizontally above a table-like display surface, the body is an operator’s hand, but the embodiment is not so limited. The pushback interaction is double-sided, so that there is an upper dead zone threshold and a lower dead zone threshold. Additionally, the linear space accessed by the pushback maneuver is provided with discrete spatial detents (e.g., “1st detent”, “2nd detent”, “3rd detent”, “4th detent”) in the upper active zone, and discrete spatial detents (e.g., “1st detent”, “2nd detent”, “3rd detent”, “4th detent”) in the lower active zone. The interaction space of an embodiment is configured so that a relatively small dead zone comprising an upper dead zone and a lower dead zone is centered at the vertical (y-axis) position at which pushback is engaged, with an active zone above the dead zone and an active zone below the dead zone.

The operator is working with an example dataset that has been analyzed into a stack of discrete parallel planes that are the data frames. The dataset may be arranged that way as a natural consequence of the physical reality it represents (e.g. discrete slices from a tomographic scan, the multiple layers of a three-dimensional integrated circuit, etc.) or because it is logical or informative to separate and discretize the data (e.g., satellite imagery acquired in a number of spectral bands, geographically organized census data with each decade’s data in a separate layer, etc.). The visual representation of the data may further be static or include dynamic elements.

During intervals when pushback functionality is not engaged, a single layer is considered ‘current’ and is represented with visual prominence by the display, and is perceived to be physically coincident with the display. Layers above and below the current layer are in this example not visually manifest (although a compact iconography is used to indicate their presence).

The operator extends his closed right hand over the display; when he opens the hand—fingers extended forward, thumb to the left, and palm pointed downward (transition: [^^^>:vx into |||:vx])—the pushback system is engaged. During a brief interval (e.g., 200 milliseconds), some number of layers adjacent to the current layer fade up with differential visibility; each is composited below or above with a blur filter and a transparency whose ‘seventies’ are dependent on the layer’s ordinal distance from the current layer.

For example, a layer (e.g., data frame) adjacent to the current layer (e.g., data frame) fades up with differential visibility as the pushback system is engaged. In this example, the stack comprises numerous data frames (any number as appropriate to datasets of the data frames) that can be traversed using the pushback system.

Simultaneously, the concentric feedback glyphs familiar from the previous example appear; in this case, the interaction

is configured so that a small dead zone is centered at the vertical (y-axis) position at which pushback is engaged, with an active zone both above and below the dead zone. This arrangement provides assistance in ‘regaining’ the original layer. The glyphs are in this case accompanied by an additional, simple graphic that indicates directed proximity to successive layers.

While the operator’s hand remains in the dead zone, no displacement of the layer stack occurs. The glyphs exhibit a ‘preparatory’ behavior identical to that in the preceding example, with the inner glyph growing as the hand nears either boundary of the zone (of course, here the behavior is double-sided and symmetric: the inner glyph is at a minimum scale at the hand’s starting y-position and grows toward coincidence with the outer glyph whether the hand moves up or down).

As the operator’s hand moves upward past the dead zone’s upper plane, the inner glyph engages the outer glyph and, as before, further movement of the hand in that direction causes anticlockwise rotational motion of the inner glyph. At the same time, the layer stack begins to ‘translate upward’: those layers above the originally-current layer take on greater transparency and blur; the originally-current layer itself becomes more transparent and more blurred; and the layers below it move toward more visibility and less blur.

In another example of upward translation of the stack, the previously-current layer takes on greater transparency (becomes invisible in this example), while the layer adjacent to the previously-current layer becomes visible as the presently-current layer. Additionally, layer adjacent to the presently-current layer fades up with differential visibility as the stack translates upward. As described above, the stack comprises numerous data frames (any number as appropriate to datasets of the data frames) that can be traversed using the pushback system.

The layer stack is configured with a mapping between real-world distances (i.e. the displacement of the operator’s hand from its initial position, as measured in room coordinates) and the ‘logical’ distance between successive layers. The translation of the layer stack is, of course, the result of this mapping, as is the instantaneous appearance of the proximity graphic, which meanwhile indicates (at first) a growing distance between the display plane and the current layer; it also indicates that the display plane is at present below the current layer.

The hand’s motion continues and the layer stack eventually passes the position at which the current layer and the next one below exactly straddle (i.e. are equidistant from) the display plane; just past this point the proximity graphic changes to indicate that the display plane is now higher than the current layer: ‘current layer status’ has now been assigned to the next lower layer. In general, the current layer is always the one closest to the physical display plane, and is the one that will be ‘selected’ when the operator disengages the pushback system.

As the operator continues to raise his hand, each consecutive layer is brought toward the display plane, becoming progressively more resolved, gaining momentary coincidence with the display plane, and then returning toward transparency and blur in favor of the next lower layer. When the operator reverses the direction of his hand’s motion, lowering it, the process is reversed, and the inner glyph rotates clockwise. As the hand eventually passes through the dead zone the stack halts with the originally-current layer in precise y-alignment with the display plane; and then y-travel of the stack resumes, bringing into successive focus those planes above

the originally-current layer. The operator's overall perception is strongly and simply that he is using his hand to push down and pull up a stack of layers.

When at last the operator releases pushback by closing his hand (or otherwise changing its pose) the system 'springs' the stack into detented y-axis alignment with the display plane, leaving as the current layer whichever was closest to the display plane as pushback was exited. During the brief interval of this positional realignment, all other layers fade back to complete transparency and the feedback glyphs smoothly vanish.

The discretized elements of the dataset (here, layers) of this example are distributed along the principal pushback (depth) axis; previously, the elements (data frames) were coplanar and arrayed laterally, along a dimension orthogonal to the depth axis. This present arrangement, along with the deployment of transparency techniques, means that data is often superimposed—some layers are viewed through others. The operator in this example nevertheless also enjoys (1) a facility for rapidly gaining neighborhood context (what are the contents of the layers above and below the current layer?); and (2) a facility for efficiently selecting and switching among parallel, stacked elements in the dataset. When the operator intends (1) alone, the provision of a dead zone allows him to return confidently to the originally selected layer. Throughout the manipulation, the suppression of two translational dimensions enables speed and accuracy (it is comparatively difficult for most humans to translate a hand vertically with no lateral drift, but the modality as described simply ignores any such lateral displacement).

It is noted that for certain purposes it may be convenient to configure the pushback input space so that the dead zone is of infinitesimal extent; then, as soon as pushback is engaged, its active mechanisms are also engaged. In the second example presented herein this would mean that the originally-current layer is treated no differently—once the pushback maneuver has begun—from any other. Empirically, the linear extent of the dead zone is a matter of operator preference.

The modalities described in this second example are pertinent across a wide variety of displays, including both two-dimensional (whether projected or emissive) and three-dimensional (whether autostereoscopic or not, aerial-image-producing or not, etc.) devices. In high-quality implementations of the latter—i.e. 3D—case, certain characteristics of the medium can vastly aid the perceptual mechanisms that underlie pushback. For example, a combination of parallax, optical depth of field, and ocular accommodation phenomena can allow multiple layers to be apprehended simultaneously, thus eliminating the need to severely fade and blur (or indeed to exclude altogether) layers distant from the display plane. The modalities apply, further, irrespective of the orientation of the display: it may be principally horizontal, as in the example, or may just as usefully be mounted at eye-height on a wall.

An extension to the scenario of this second example depicts the usefulness of two-handed manipulation. In certain applications, translating either the entire layer stack or an individual layer laterally (i.e. in the x and z directions) is necessary. In an embodiment, the operator's other—that is, non-pushback—hand can effect this transformation, for example through a modality in which bringing the hand into close proximity to the display surface allows one of the dataset's layers to be 'slid around', so that its offset x-z position follows that of the hand.

Operators may generally find it convenient and easily tractable to undertake lateral translation and pushback manipulations simultaneously. It is perhaps not wholly fatuous to

propose that the assignment of continuous-domain manipulations to one hand and discrete-style work to the other may act to optimize cognitive load.

It is informative to consider yet another example of pushback under the SOE in which there is no natural visual aspect to the dataset. Representative is the problem of monitoring a plurality of audio channels and of intermittently selecting one from among the collection. An application of the pushback system enables such a task in an environment outfitted for aural but not visual output; the modality is remarkably similar to that of the preceding example.

An operator, standing or seated, is listening to a single channel of audio. Conceptually, this audio exists in the vertical plane—called the 'aural plane'—that geometrically includes her ears; additional channels of audio are resident in additional planes parallel to the aural plane but displaced forward and back, along the z-axis.

Opening her hand, held nine inches in front of her, with palm facing forward, she engages the pushback system. The audio in several proximal planes fades up differentially; the volume of each depends inversely on its ordinal distance from the current channel's plane. In practice, it is perceptually unrealistic to allow more than two or four additional channels to become audible. At the same time, an 'audio glyph' fades up to provide proximity feedback. Initially, while the operator's hand is held in the dead zone, the glyph is a barely audible two-note chord (initially in unison).

As the operator moves her hand forward or backward through the dead zone, the volumes of the audio channels remain fixed while that of the glyph increases. When the hand crosses the front or rear threshold of the dead zone, the glyph reaches its 'active' volume (which is still subordinate to the current channel's volume).

Once the operator's hand begins moving through the active zone—in the forward direction, say—the expected effect on the audio channels obtains: the current channel plane is pushed farther from the aural plane, and its volume (and the volumes of those channels still farther forward) is progressively reduced. The volume of each 'dorsal' channel plane, on the other hand, increases as it nears the aural plane.

The audio glyph, meanwhile, has switched modes. The hand's forward progress is accompanied by the rise in frequency of one of the tones; at the 'midway point', when the aural plane bisects one audio channel plane and the next, the tones form an exact fifth (mathematically, it should be a tritone interval, but there is an abundance of reasons that this is to be eschewed). The variable tone's frequency continues rising as the hand continues farther forward, until eventually the operator 'reaches' the next audio plane, at which point the tones span precisely an octave.

Audition of the various channels proceeds, the operator translating her hand forward and back to access each in turn. Finally, to select one she merely closes her hand, concluding the pushback session and causing the collection of audio planes to 'spring' into alignment. The other (non-selected) channels fade to inaudibility, as does the glyph.

This example has illustrated a variant on pushback application in which the same facilities are again afforded: access to neighborhood context and rapid selection of discretized data element (here, an individual audio stream). The scenario substitutes an aural feedback mechanism, and in particular one that exploits the reliable human capacity for discerning certain frequency intervals, to provide the operator with information about whether she is 'close enough' to a target channel to make a selection. This is particularly important in the case of voice channels, in which 'audible' signals are only inter-

mittently present; the continuous nature of the audio feedback glyph leaves it present and legible even when the channel itself has gone silent.

It is noted that if the SOE in this present example includes the capacity for spatialized audio, the perception of successive audio layers receding into the forward distance and approaching from the back (or vice versa) may be greatly enhanced. Further, the opportunity to more literally ‘locate’ the selected audio plane at the position of the operator, with succeeding layers in front of the operator and preceding layers behind, is usefully exploitable.

Other instantiations of the audio glyph are possible, and indeed the nature of the various channels’ contents, including their spectral distributions, tends to dictate which kind of glyph will be most clearly discernible. By way of example, another audio glyph format maintains constant volume but employs periodic clicking, with the interval between clicks proportional to the proximity between the aural plane and the closest audio channel plane. Finally, under certain circumstances, and depending on the acuity of the operator, it is possible to use audio pushback with no feedback glyph at all.

With reference to the pushback mechanism, as the number and density of spatial detents in the dataset’s representation increases toward the very large, the space and its parameterization becomes effectively continuous—that is to say, non-detented. Pushback remains nonetheless effective at such extremes, in part because the dataset’s ‘initial state’ prior to each invocation of pushback may be treated as a temporary detent, realized simply as a dead zone.

An application of such non-detented pushback may be found in connection with the idea of an infinitely (or at least substantially) zoomable diagram. Pushback control of zoom functionality associates offset hand position with affine scale value, so that as the operator pushes his hand forward or back the degree of zoom decreases or increases (respectively). The original, pre-pushback zoom state is always readily accessible, however, because the direct mapping of position to zoom parameter insures that returning the control hand to the dead zone also effects return of the zoom value to its initial state.

Each scenario described in the examples above provides a description of the salient aspects of the pushback system and its use under the SOE. It should further be understood that each of the maneuvers described herein can be accurately and comprehensibly undertaken in a second or less, because of the efficiency and precision enabled by allowing a particular kind of perceptual feedback to guide human movement. At other times, operators also find it useful to remain in a single continuous pushback ‘session’ for tens of seconds: exploratory and context-acquisition goals are well served by pushback over longer intervals.

The examples described above employed a linear mapping of physical input (gesture) space to representational space: translating the control hand by A units in real space always results in a translation by B units [prime] in the representational space, irrespective of the real-space position at which the A-translation is undertaken. However, other mappings are possible. In particular, the degree of fine motor control enjoyed by most human operators allows the use of nonlinear mappings, in which for example differential gestural translations far from the active threshold can translate into larger displacements along the parameterized dimension than do gestural translations near the threshold.

Coincident Virtual/Display and Physical Spaces

The system can provide an environment in which virtual space depicted on one or more display devices (“screens”) is treated as coincident with the physical space inhabited by the

operator or operators of the system. An embodiment of such an environment is described here. This current embodiment includes three projector-driven screens at fixed locations, is driven by a single desktop computer, and is controlled using the gestural vocabulary and interface system described herein. Note, however, that any number of screens are supported by the techniques being described; that those screens may be mobile (rather than fixed); that the screens may be driven by many independent computers simultaneously; and that the overall system can be controlled by any input device or technique.

The interface system described in this disclosure should have a means of determining the dimensions, orientations and positions of screens in physical space. Given this information, the system is able to dynamically map the physical space in which these screens are located (and which the operators of the system inhabit) as a projection into the virtual space of computer applications running on the system. As part of this automatic mapping, the system also translates the scale, angles, depth, dimensions and other spatial characteristics of the two spaces in a variety of ways, according to the needs of the applications that are hosted by the system.

This continuous translation between physical and virtual space makes possible the consistent and pervasive use of a number of interface techniques that are difficult to achieve on existing application platforms or that must be implemented piece-meal for each application running on existing platforms. These techniques include (but are not limited to):

1) Use of “literal pointing”—using the hands in a gestural interface environment, or using physical pointing tools or devices—as a pervasive and natural interface technique.

2) Automatic compensation for movement or repositioning of screens.

3) Graphics rendering that changes depending on operator position, for example simulating parallax shifts to enhance depth perception.

4) Inclusion of physical objects in on-screen display—taking into account real-world position, orientation, state, etc. For example, an operator standing in front of a large, opaque screen, could see both applications graphics and a representation of the true position of a scale model that is behind the screen (and is, perhaps, moving or changing orientation).

It is important to note that literal pointing is different from the abstract pointing used in mouse-based windowing interfaces and most other contemporary systems. In those systems, the operator must learn to manage a translation between a virtual pointer and a physical pointing device, and must map between the two cognitively.

By contrast, in the systems described in this disclosure, there is no difference between virtual and physical space (except that virtual space is more amenable to mathematical manipulation), either from an application or user perspective, so there is no cognitive translation required of the operator.

The closest analogy for the literal pointing provided by the embodiment described here is the touch-sensitive screen (as found, for example, on many ATM machines). A touch-sensitive screen provides a one to one mapping between the two-dimensional display space on the screen and the two-dimensional input space of the screen surface. In an analogous fashion, the systems described here provide a flexible mapping (possibly, but not necessarily, one to one) between a virtual space displayed on one or more screens and the physical space inhabited by the operator. Despite the usefulness of the analogy, it is worth understanding that the extension of this “mapping approach” to three dimensions, an arbitrarily large architectural environment, and multiple screens is non-trivial.

In addition to the components described herein, the system may also implement algorithms implementing a continuous, systems-level mapping (perhaps modified by rotation, translation, scaling or other geometrical transformations) between the physical space of the environment and the display space on each screen.

A rendering stack which takes the computational objects and the mapping and outputs a graphical representation of the virtual space.

An input events processing stack which takes event data from a control system (in the current embodiment both gestural and pointing data from the system and mouse input) and maps spatial data from input events to coordinates in virtual space. Translated events are then delivered to running applications.

A “glue layer” allowing the system to host applications running across several computers on a local area network.

Embodiments of a spatial-continuum input system are described herein as comprising network-based data representation, transit, and interchange that includes a system called “plasma” that comprises subsystems “slawx”, “proteins”, and “pools”, as described in detail below. The pools and proteins are components of methods and systems described herein for encapsulating data that is to be shared between or across processes. These mechanisms also include slawx (plural of “slaw”) in addition to the proteins and pools. Generally, slawx provide the lowest-level of data definition for inter-process exchange, proteins provide mid-level structure and hooks for querying and filtering, and pools provide for high-level organization and access semantics. Slawx include a mechanism for efficient, platform-independent data representation and access. Proteins provide a data encapsulation and transport scheme using slawx as the payload. Pools provide structured and flexible aggregation, ordering, filtering, and distribution of proteins within a process, among local processes, across a network between remote or distributed processes, and via longer term (e.g. on-disk, etc.) storage.

The configuration and implementation of the embodiments described herein include several constructs that together enable numerous capabilities. For example, the embodiments described herein provide efficient exchange of data between large numbers of processes as described above. The embodiments described herein also provide flexible data “typing” and structure, so that widely varying kinds and uses of data are supported. Furthermore, embodiments described herein include flexible mechanisms for data exchange (e.g., local memory, disk, network, etc.), all driven by substantially similar application programming interfaces (APIs). Moreover, embodiments described herein enable data exchange between processes written in different programming languages. Additionally, embodiments described herein enable automatic maintenance of data caching and aggregate state.

FIG. 14 is a block diagram of a processing environment including data representations using slawx, proteins, and pools, under an embodiment. The principal constructs of the embodiments presented herein include slawx (plural of “slaw”), proteins, and pools. Slawx as described herein includes a mechanism for efficient, platform-independent data representation and access. Proteins, as described in detail herein, provide a data encapsulation and transport scheme, and the payload of a protein of an embodiment includes slawx. Pools, as described herein, provide structured yet flexible aggregation, ordering, filtering, and distribution of proteins. The pools provide access to data, by virtue of proteins, within a process, among local processes, across a network between remote or distributed processes, and via ‘longer term’ (e.g. on-disk) storage.

FIG. 15 is a block diagram of a protein, under an embodiment. The protein includes a length header, a descrip, and an ingest. Each of the descrip and ingest includes slaw or slawx, as described in detail below.

FIG. 16 is a block diagram of a descrip, under an embodiment. The descrip includes an offset, a length, and slawx, as described in detail below.

FIG. 17 is a block diagram of an ingest, under an embodiment. The ingest includes an offset, a length, and slawx, as described in detail below.

FIG. 18 is a block diagram of a slaw, under an embodiment. The slaw includes a type header and type-specific data, as described in detail below.

FIG. 19A is a block diagram of a protein in a pool, under an embodiment. The protein includes a length header (“protein length”), a descripts offset, an ingests offset, a descrip, and an ingest. The descripts includes an offset, a length, and a slaw. The ingest includes an offset, a length, and a slaw.

The protein as described herein is a mechanism for encapsulating data that needs to be shared between processes, or moved across a bus or network or other processing structure. As an example, proteins provide an improved mechanism for transport and manipulation of data including data corresponding to or associated with user interface events; in particular, the user interface events of an embodiment include those of the gestural interface described above. As a further example, proteins provide an improved mechanism for transport and manipulation of data including, but not limited to, graphics data or events, and state information, to name a few. A protein is a structured record format and an associated set of methods for manipulating records. Manipulation of records as used herein includes putting data into a structure, taking data out of a structure, and querying the format and existence of data. Proteins are configured to be used via code written in a variety of computer languages. Proteins are also configured to be the basic building block for pools, as described herein. Furthermore, proteins are configured to be natively able to move between processors and across networks while maintaining intact the data they include.

In contrast to conventional data transport mechanisms, proteins are untyped. While being untyped, the proteins provide a powerful and flexible pattern-matching facility, on top of which “type-like” functionality is implemented. Proteins configured as described herein are also inherently multi-point (although point-to-point forms are easily implemented as a subset of multi-point transmission). Additionally, proteins define a “universal” record format that does not differ (or differs only in the types of optional optimizations that are performed) between in-memory, on-disk, and on-the-wire (network) formats, for example.

Referring to FIGS. 15 and 19A, a protein of an embodiment is a linear sequence of bytes. Within these bytes are encapsulated a descripts list and a set of key-value pairs called ingests. The descripts list includes an arbitrarily elaborate but efficiently filterable per-protein event description. The ingests include a set of key-value pairs that comprise the actual contents of the protein.

Proteins’ concern with key-value pairs, as well as some core ideas about network-friendly and multi-point data interchange, is shared with earlier systems that privilege the concept of “tuples” (e.g., Linda, Jini). Proteins differ from tuple-oriented systems in several major ways, including the use of the descripts list to provide a standard, optimizable pattern matching substrate. Proteins also differ from tuple-oriented systems in the rigorous specification of a record format

appropriate for a variety of storage and language constructs, along with several particular implementations of “interfaces” to that record format.

Turning to a description of proteins, the first four or eight bytes of a protein specify the protein’s length, which must be a multiple of 16 bytes in an embodiment. This 16-byte granularity ensures that byte-alignment and bus-alignment efficiencies are achievable on contemporary hardware. A protein that is not naturally “quad-word aligned” is padded with arbitrary bytes so that its length is a multiple of 16 bytes.

The length portion of a protein has the following format: 32 bits specifying length, in big-endian format, with the four lowest-order bits serving as flags to indicate macro-level protein structure characteristics; followed by 32 further bits if the protein’s length is greater than 2¹² bytes.

The 16-byte-alignment proviso of an embodiment means that the lowest order bits of the first four bytes are available as flags. And so the first three low-order bit flags indicate whether the protein’s length can be expressed in the first four bytes or requires eight, whether the protein uses big-endian or little-endian byte ordering, and whether the protein employs standard or non-standard structure, respectively, but the protein is not so limited. The fourth flag bit is reserved for future use.

If the eight-byte length flag bit is set, the length of the protein is calculated by reading the next four bytes and using them as the high-order bytes of a big-endian, eight-byte integer (with the four bytes already read supplying the low-order portion). If the little-endian flag is set, all binary numerical data in the protein is to be interpreted as little-endian (otherwise, big-endian). If the non-standard flag bit is set, the remainder of the protein does not conform to the standard structure to be described below.

Non-standard protein structures will not be discussed further herein, except to say that there are various methods for describing and synchronizing on non-standard protein formats available to a systems programmer using proteins and pools, and that these methods can be useful when space or compute cycles are constrained. For example, the shortest protein of an embodiment is sixteen bytes. A standard-format protein cannot fit any actual payload data into those sixteen bytes (the lion’s share of which is already relegated to describing the location of the protein’s component parts). But a non-standard format protein could conceivably use 12 of its 16 bytes for data. Two applications exchanging proteins could mutually decide that any 16-byte-long proteins that they emit always include 12 bytes representing, for example, 12 8-bit sensor values from a real-time analog-to-digital converter.

Immediately following the length header, in the standard structure of a protein, two more variable-length integer numbers appear. These numbers specify offsets to, respectively, the first element in the descripts list and the first key-value pair (ingest). These offsets are also referred to herein as the descripts offset and the ingests offset, respectively. The byte order of each quad of these numbers is specified by the protein endianness flag bit. For each, the most significant bit of the first four bytes determines whether the number is four or eight bytes wide. If the most significant bit (msb) is set, the first four bytes are the most significant bytes of a double-word (eight byte) number. This is referred to herein as “offset form”. Use of separate offsets pointing to descripts and pairs allows descripts and pairs to be handled by different code paths, making possible particular optimizations relating to, for example, descripts pattern-matching and protein assembly. The presence of these two offsets at the beginning of a protein also allows for several useful optimizations.

Most proteins will not be so large as to require eight-byte lengths or pointers, so in general the length (with flags) and two offset numbers will occupy only the first three bytes of a protein. On many hardware or system architectures, a fetch or read of a certain number of bytes beyond the first is “free” (e.g., 16 bytes take exactly the same number of clock cycles to pull across the Cell processor’s main bus as a single byte).

In many instances it is useful to allow implementation-specific or context-specific caching or metadata inside a protein. The use of offsets allows for a “hole” of arbitrary size to be created near the beginning of the protein, into which such metadata may be slotted. An implementation that can make use of eight bytes of metadata gets those bytes for free on many system architectures with every fetch of the length header for a protein.

The descripts offset specifies the number of bytes between the beginning of the protein and the first descrip entry. Each descrip entry comprises an offset (in offset form, of course) to the next descrip entry, followed by a variable-width length field (again in offset format), followed by a slaw. If there are no further descripts, the offset is, by rule, four bytes of zeros. Otherwise, the offset specifies the number of bytes between the beginning of this descrip entry and a subsequent descrip entry. The length field specifies the length of the slaw, in bytes.

In most proteins, each descrip is a string, formatted in the slaw string fashion: a four-byte length/type header with the most significant bit set and only the lower 30 bits used to specify length, followed by the header’s indicated number of data bytes. As usual, the length header takes its endianness from the protein. Bytes are assumed to encode UTF-8 characters (and thus—not a bene—the number of characters is not necessarily the same as the number of bytes).

The ingests offset specifies the number of bytes between the beginning of the protein and the first ingest entry. Each ingest entry comprises an offset (in offset form) to the next ingest entry, followed again by a length field and a slaw. The ingests offset is functionally identical to the descripts offset, except that it points to the next ingest entry rather than to the next descrip entry.

In most proteins, every ingest is of the slaw cons type comprising a two-value list, generally used as a key/value pair. The slaw cons record comprises a four-byte length/type header with the second most significant bit set and only the lower 30 bits used to specify length; a four-byte offset to the start of the value (second) element; the four-byte length of the key element; the slaw record for the key element; the four-byte length of the value element; and finally the slaw record for the value element.

Generally, the cons key is a slaw string. The duplication of data across the several protein and slaw cons length and offsets field provides yet more opportunity for refinement and optimization.

The construct used under an embodiment to embed typed data inside proteins, as described above, is a tagged byte-sequence specification and abstraction called a “slaw” (the plural is “slawx”). A slaw is a linear sequence of bytes representing a piece of (possibly aggregate) typed data, and is associated with programming-language-specific APIs that allow slawx to be created, modified and moved around between memory spaces, storage media, and machines. The slaw type scheme is intended to be extensible and as lightweight as possible, and to be a common substrate that can be used from any programming language.

The desire to build an efficient, large-scale inter-process communication mechanism is the driver of the slaw configuration. Conventional programming languages provide

sophisticated data structures and type facilities that work well in process-specific memory layouts, but these data representations invariably break down when data needs to be moved between processes or stored on disk. The slaw architecture is, first, a substantially efficient, multi-platform friendly, low-level data model for inter-process communication.

But even more importantly, slawx are configured to influence, together with proteins, and enable the development of future computing hardware (microprocessors, memory controllers, disk controllers). A few specific additions to, say, the instruction sets of commonly available microprocessors make it possible for slawx to become as efficient even for single-process, in-memory data layout as the schema used in most programming languages.

Each slaw comprises a variable-length type header followed by a type-specific data layout. In an example embodiment, which supports full slaw functionality in C, C++ and Ruby for example, types are indicated by a universal integer defined in system header files accessible from each language. More sophisticated and flexible type resolution functionality is also enabled: for example, indirect typing via universal object IDs and network lookup.

The slaw configuration of an embodiment allows slaw records to be used as objects in language-friendly fashion from both Ruby and C++, for example. A suite of utilities external to the C++ compiler sanity-check slaw byte layout, create header files and macros specific to individual slaw types, and auto-generate bindings for Ruby. As a result, well-configured slaw types are quite efficient even when used from within a single process. Any slaw anywhere in a process's accessible memory can be addressed without a copy or "deserialization" step.

Slaw functionality of an embodiment includes API facilities to perform one or more of the following: create a new slaw of a specific type; create or build a language-specific reference to a slaw from bytes on disk or in memory; embed data within a slaw in type-specific fashion; query the size of a slaw; retrieve data from within a slaw; clone a slaw; and translate the endianness and other format attributes of all data within a slaw. Every species of slaw implements the above behaviors.

FIGS. 19B/1 and 19B2 show a slaw header format, under an embodiment. A detailed description of the slaw follows.

The internal structure of each slaw optimizes each of type resolution, access to encapsulated data, and size information for that slaw instance. In an embodiment, the full set of slaw types is by design minimally complete, and includes: the slaw string; the slaw cons (i.e. dyad); the slaw list; and the slaw numerical object, which itself represents a broad set of individual numerical types understood as permutations of a half-dozen or so basic attributes. The other basic property of any slaw is its size. In an embodiment, slawx have byte-lengths quantized to multiples of four; these four-byte words are referred to herein as 'quads'. In general, such quad-based sizing aligns slawx well with the configurations of modern computer hardware architectures.

The first four bytes of every slaw in an embodiment comprise a header structure that encodes type-description and other meta-information, and that ascribes specific type meanings to particular bit patterns. For example, the first (most significant) bit of a slaw header is used to specify whether the size (length in quad-words) of that slaw follows the initial four-byte type header. When this bit is set, it is understood that the size of the slaw is explicitly recorded in the next four bytes of the slaw (e.g., bytes five through eight); if the size of the slaw is such that it cannot be represented in four bytes (i.e. if the size is or is larger than two to the thirty-second power) then the next-most-significant bit of the slaw's initial four

bytes is also set, which means that the slaw has an eight-byte (rather than four byte) length. In that case, an inspecting process will find the slaw's length stored in ordinal bytes five through twelve. On the other hand, the small number of slaw types means that in many cases a fully specified typical bit-pattern "leaves unused" many bits in the four byte slaw header; and in such cases these bits may be employed to encode the slaw's length, saving the bytes (five through eight) that would otherwise be required.

For example, an embodiment leaves the most significant bit of the slaw header (the "length follows" flag) unset and sets the next bit to indicate that the slaw is a "wee cons", and in this case the length of the slaw (in quads) is encoded in the remaining thirty bits. Similarly, a "wee string" is marked by the pattern 001 in the header, which leaves twenty-nine bits for representation of the slaw-string's length; and a leading 0001 in the header describes a "wee list", which by virtue of the twenty-eight available length-representing bits can be a slaw list of up to two-to-the-twenty-eight quads in size. A "full string" (or cons or list) has a different bit signature in the header, with the most significant header bit necessarily set because the slaw length is encoded separately in bytes five through eight (or twelve, in extreme cases). Note that the Plasma implementation "decides" at the instant of slaw construction whether to employ the "wee" or the "full" version of these constructs (the decision is based on whether the resulting size will "fit" in the available wee bits or not), but the full-vs.-wee detail is hidden from the user of the Plasma implementation, who knows and cares only that she is using a slaw string, or a slaw cons, or a slaw list.

Numeric slawx are, in an embodiment, indicated by the leading header pattern 00001. Subsequent header bits are used to represent a set of orthogonal properties that may be combined in arbitrary permutation. An embodiment employs, but is not limited to, five such character bits to indicate whether or not the number is: (1) floating point; (2) complex; (3) unsigned; (4) "wide"; (5) "stumpy" ((4) "wide" and (5) "stumpy" are permuted to indicate eight, sixteen, thirty-two, and sixty-four bit number representations). Two additional bits (e.g., (7) and (8)) indicate that the encapsulated numeric data is a two-, three-, or four-element vector (with both bits being zero suggesting that the numeric is a "one-element vector" (i.e. a scalar)). In this embodiment the eight bits of the fourth header byte are used to encode the size (in bytes, not quads) of the encapsulated numeric data. This size encoding is offset by one, so that it can represent any size between and including one and two hundred fifty-six bytes. Finally, two character bits (e.g., (9) and (10)) are used to indicate that the numeric data encodes an array of individual numeric entities, each of which is of the type described by character bits (1) through (8). In the case of an array, the individual numeric entities are not each tagged with additional headers, but are packed as continuous data following the single header and, possibly, explicit slaw size information.

This embodiment affords simple and efficient slaw duplication (which can be implemented as a byte-for-byte copy) and extremely straightforward and efficient slaw comparison (two slawx are the same in this embodiment if and only if there is a one-to-one match of each of their component bytes considered in sequence). This latter property is important, for example, to an efficient implementation of the protein architecture, one of whose critical and pervasive features is the ability to search through or 'match on' a protein's descripts list.

Further, the embodiments herein allow aggregate slaw forms (e.g., the slaw cons and the slaw list) to be constructed simply and efficiently. For example, an embodiment builds a

slaw cons from two component slawx, which may be of any type, including themselves aggregates, by: (a) querying each component slaw's size; (b) allocating memory of size equal to the sum of the sizes of the two component slawx and the one, two, or three quads needed for the header-plus-size structure; (c) recording the slaw header (plus size information) in the first four, eight, or twelve bytes; and then (d) copying the component slawx's bytes in turn into the immediately succeeding memory. Significantly, such a construction routine need know nothing about the types of the two component slawx; only their sizes (and accessibility as a sequence of bytes) matters. The same process pertains to the construction of slaw lists, which are ordered encapsulations of arbitrarily many sub-slawx of (possibly) heterogeneous type.

A further consequence of the slaw system's fundamental format as sequential bytes in memory obtains in connection with "traversal" activities—a recurring use pattern uses, for example, sequential access to the individual slawx stored in a slaw list. The individual slawx that represent the descripts and ingests within a protein structure must similarly be traversed. Such maneuvers are accomplished in a stunningly straightforward and efficient manner: to "get to" the next slaw in a slaw list, one adds the length of the current slaw to its location in memory, and the resulting memory location is identically the header of the next slaw. Such simplicity is possible because the slaw and protein design eschews "indirection"; there are no pointers; rather, the data simply exists, in its totality, in situ.

To the point of slaw comparison, a complete implementation of the Plasma system must acknowledge the existence of differing and incompatible data representation schemes across and among different operating systems, CPUs, and hardware architectures. Major such differences include byte-ordering policies (e.g., little- vs. big-endianness) and floating-point representations; other differences exist. The Plasma specification requires that the data encapsulated by slawx be guaranteed interpretable (i.e., must appear in the native format of the architecture or platform from which the slaw is being inspected. This requirement means in turn that the Plasma system is itself responsible for data format conversion. However, the specification stipulates only that the conversion take place before a slaw becomes "at all visible" to an executing process that might inspect it. It is therefore up to the individual implementation at which point it chooses to perform such format conversion; two appropriate approaches are that slaw data payloads are conformed to the local architecture's data format (1) as an individual slaw is "pulled out" of a protein in which it had been packed, or (2) for all slaw in a protein simultaneously, as that protein is extracted from the pool in which it was resident. Note that the conversion stipulation considers the possibility of hardware-assisted implementations. For example, networking chipsets built with explicit Plasma capability may choose to perform format conversion intelligently and at the "instant of transmission", based on the known characteristics of the receiving system. Alternately, the process of transmission may convert data payloads into a canonical format, with the receiving process symmetrically converting from canonical to "local" format. Another embodiment performs format conversion "at the metal", meaning that data is always stored in canonical format, even in local memory, and that the memory controller hardware itself performs the conversion as data is retrieved from memory and placed in the registers of the proximal CPU.

A minimal (and read-only) protein implementation of an embodiment includes operation or behavior in one or more applications or programming languages making use of pro-

teins. FIG. 19C is a flow diagram 650 for using proteins, under an embodiment. Operation begins by querying 652 the length in bytes of a protein. The number of descripts entries is queried 654. The number of ingests is queried 656. A descript entry is retrieved 658 by index number. An ingest is retrieved 660 by index number.

The embodiments described herein also define basic methods allowing proteins to be constructed and filled with data, helper-methods that make common tasks easier for programmers, and hooks for creating optimizations. FIG. 19D is a flow diagram 670 for constructing or generating proteins, under an embodiment. Operation begins with creation 672 of a new protein. A series of descripts entries are appended 674. An ingest is also appended 676. The presence of a matching descript is queried 678, and the presence of a matching ingest key is queried 680. Given an ingest key, an ingest value is retrieved 682. Pattern matching is performed 684 across descripts. Non-structured metadata is embedded 686 near the beginning of the protein.

As described above, slawx provide the lowest-level of data definition for inter-process exchange, proteins provide mid-level structure and hooks for querying and filtering, and pools provide for high-level organization and access semantics. The pool is a repository for proteins, providing linear sequencing and state caching. The pool also provides multi-process access by multiple programs or applications of numerous different types. Moreover, the pool provides a set of common, optimizable filtering and pattern-matching behaviors.

The pools of an embodiment, which can accommodate tens of thousands of proteins, function to maintain state, so that individual processes can offload much of the tedious book-keeping common to multi-process program code. A pool maintains or keeps a large buffer of past proteins available—the Platonic pool is explicitly infinite—so that participating processes can scan both backwards and forwards in a pool at will. The size of the buffer is implementation dependent, of course, but in common usage it is often possible to keep proteins in a pool for hours or days.

The most common style of pool usage as described herein hews to a biological metaphor, in contrast to the mechanistic, point-to-point approach taken by existing inter-process communication frameworks. The name protein alludes to biological inspiration: data proteins in pools are available for flexible querying and pattern matching by a large number of computational processes, as chemical proteins in a living organism are available for pattern matching and filtering by large numbers of cellular agents.

Two additional abstractions lean on the biological metaphor, including use of "handlers", and the Golgi framework. A process that participates in a pool generally creates a number of handlers. Handlers are relatively small bundles of code that associate match conditions with handle behaviors. By tying one or more handlers to a pool, a process sets up flexible call-back triggers that encapsulate state and react to new proteins.

A process that participates in several pools generally inherits from an abstract Golgi class. The Golgi framework provides a number of useful routines for managing multiple pools and handlers. The Golgi class also encapsulates parent-child relationships, providing a mechanism for local protein exchange that does not use a pool.

A pools API provided under an embodiment is configured to allow pools to be implemented in a variety of ways, in order to account both for system-specific goals and for the available capabilities of given hardware and network architectures. The two fundamental system provisions upon which pools depend are a storage facility and a means of inter-process communi-

cation. The extant systems described herein use a flexible combination of shared memory, virtual memory, and disk for the storage facility, and IPC queues and TCP/IP sockets for inter-process communication.

Pool functionality of an embodiment includes, but is not limited to, the following: participating in a pool; placing a protein in a pool; retrieving the next unseen protein from a pool; rewinding or fast-forwarding through the contents (e.g., proteins) within a pool. Additionally, pool functionality can include, but is not limited to, the following: setting up a streaming pool call-back for a process; selectively retrieving proteins that match particular patterns of descripts or ingests keys; scanning backward and forwards for proteins that match particular patterns of descripts or ingests keys.

The proteins described above are provided to pools as a way of sharing the protein data contents with other applications. FIG. 20 is a block diagram of a processing environment including data exchange using slawx, proteins, and pools, under an embodiment. This example environment includes three devices (e.g., Device X, Device Y, and Device Z, collectively referred to herein as the "devices") sharing data through the use of slawx, proteins and pools as described above. Each of the devices is coupled to the three pools (e.g., Pool 1, Pool 2, Pool 3). Pool 1 includes numerous proteins (e.g., Protein X1, Protein Z2, Protein Y2, Protein X4, Protein Y4) contributed or transferred to the pool from the respective devices (e.g., protein Z2 is transferred or contributed to pool 1 by device Z, etc.). Pool 2 includes numerous proteins (e.g., Protein Z4, Protein Y3, Protein Z1, Protein X3) contributed or transferred to the pool from the respective devices (e.g., protein Y3 is transferred or contributed to pool 2 by device Y, etc.). Pool 3 includes numerous proteins (e.g., Protein Y1, Protein Z3, Protein X2) contributed or transferred to the pool from the respective devices (e.g., protein X2 is transferred or contributed to pool 3 by device X, etc.). While the example described above includes three devices coupled or connected among three pools, any number of devices can be coupled or connected in any manner or combination among any number of pools, and any pool can include any number of proteins contributed from any number or combination of devices. The proteins and pools of this example are as described above with reference to FIGS. 18-23.

FIG. 21 is a block diagram of a processing environment including multiple devices and numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the devices, under an embodiment. This system is but one example of a multi-user, multi-device, multi-computer interactive control scenario or configuration. More particularly, in this example, an interactive system, comprising multiple devices (e.g., device A, B, etc.) and a number of programs (e.g., apps AA-AX, apps BA-BX, etc.) running on the devices uses the Plasma constructs (e.g., pools, proteins, and slaw) to allow the running programs to share and collectively respond to the events generated by these input devices.

In this example, each device (e.g., device A, B, etc.) translates discrete raw data generated by or output from the programs (e.g., apps AA-AX, apps BA-BX, etc.) running on that respective device into Plasma proteins and deposits those proteins into a Plasma pool. For example, program AX generates data or output and provides the output to device A which, in turn, translates the raw data into proteins (e.g., protein 1A, protein 2A, etc.) and deposits those proteins into the pool. As another example, program BC generates data and

provides the data to device B which, in turn, translates the data into proteins (e.g., protein 1B, protein 2B, etc.) and deposits those proteins into the pool.

Each protein contains a descrip list that specifies the data or output registered by the application as well as identifying information for the program itself. Where possible, the protein descripts may also ascribe a general semantic meaning for the output event or action. The protein's data payload (e.g., ingests) carries the full set of useful state information for the program event.

The proteins, as described above, are available in the pool for use by any program or device coupled or connected to the pool, regardless of type of the program or device. Consequently, any number of programs running on any number of computers may extract event proteins from the input pool. These devices need only be able to participate in the pool via either the local memory bus or a network connection in order to extract proteins from the pool. An immediate consequence of this is the beneficial possibility of decoupling processes that are responsible for generating processing events from those that use or interpret the events. Another consequence is the multiplexing of sources and consumers of events so that devices may be controlled by one person or may be used simultaneously by several people (e.g., a Plasma-based input framework supports many concurrent users), while the resulting event streams are in turn visible to multiple event consumers.

As an example, device C can extract one or more proteins (e.g., protein 1A, protein 2A, etc.) from the pool. Following protein extraction, device C can use the data of the protein, retrieved or read from the slaw of the descripts and ingests of the protein, in processing events to which the protein data corresponds. As another example, device B can extract one or more proteins (e.g., protein 1C, protein 2A, etc.) from the pool. Following protein extraction, device B can use the data of the protein in processing events to which the protein data corresponds.

Devices and/or programs coupled or connected to a pool may skim backwards and forwards in the pool looking for particular sequences of proteins. It is often useful, for example, to set up a program to wait for the appearance of a protein matching a certain pattern, then skim backwards to determine whether this protein has appeared in conjunction with certain others. This facility for making use of the stored event history in the input pool often makes writing state management code unnecessary, or at least significantly reduces reliance on such undesirable coding patterns.

FIG. 22 is a block diagram of a processing environment including multiple devices and numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the devices, under an alternative embodiment. This system is but one example of a multi-user, multi-device, multi-computer interactive control scenario or configuration. More particularly, in this example, an interactive system, comprising multiple devices (e.g., devices X and Y coupled to devices A and B, respectively) and a number of programs (e.g., apps AA-AX, apps BA-BX, etc.) running on one or more computers (e.g., device A, device B, etc.) uses the Plasma constructs (e.g., pools, proteins, and slaw) to allow the running programs to share and collectively respond to the events generated by these input devices.

In this example, each device (e.g., devices X and Y coupled to devices A and B, respectively) is managed and/or coupled to run under or in association with one or more programs hosted on the respective device (e.g., device A, device B, etc.)

which translates the discrete raw data generated by the device (e.g., device X, device A, device Y, device B, etc.) hardware into Plasma proteins and deposits those proteins into a Plasma pool. For example, device X running in association with application AB hosted on device A generates raw data, translates the discrete raw data into proteins (e.g., protein 1A, protein 2A, etc.) and deposits those proteins into the pool. As another example, device X running in association with application AT hosted on device A generates raw data, translates the discrete raw data into proteins (e.g., protein 1A, protein 2A, etc.) and deposits those proteins into the pool. As yet another example, device Z running in association with application CD hosted on device C generates raw data, translates the discrete raw data into proteins (e.g., protein 1C, protein 2C, etc.) and deposits those proteins into the pool.

Each protein contains a descrip list that specifies the action registered by the input device as well as identifying information for the device itself. Where possible, the protein describes may also ascribe a general semantic meaning for the device action. The protein's data payload (e.g., ingests) carries the full set of useful state information for the device event.

The proteins, as described above, are available in the pool for use by any program or device coupled or connected to the pool, regardless of type of the program or device. Consequently, any number of programs running on any number of computers may extract event proteins from the input pool. These devices need only be able to participate in the pool via either the local memory bus or a network connection in order to extract proteins from the pool. An immediate consequence of this is the beneficial possibility of decoupling processes that are responsible for generating processing events from those that use or interpret the events. Another consequence is the multiplexing of sources and consumers of events so that input devices may be controlled by one person or may be used simultaneously by several people (e.g., a Plasma-based input framework supports many concurrent users), while the resulting event streams are in turn visible to multiple event consumers.

Devices and/or programs coupled or connected to a pool may skim backwards and forwards in the pool looking for particular sequences of proteins. It is often useful, for example, to set up a program to wait for the appearance of a protein matching a certain pattern, then skim backwards to determine whether this protein has appeared in conjunction with certain others. This facility for making use of the stored event history in the input pool often makes writing state management code unnecessary, or at least significantly reduces reliance on such undesirable coding patterns.

FIG. 23 is a block diagram of a processing environment including multiple input devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the events generated by the input devices, under another alternative embodiment. This system is but one example of a multi-user, multi-device, multi-computer interactive control scenario or configuration. More particularly, in this example, an interactive system, comprising multiple input devices (e.g., input devices A, B, BA, and BB, etc.) and a number of programs (not shown) running on one or more computers (e.g., device A, device B, etc.) uses the Plasma constructs (e.g., pools, proteins, and slaw) to allow the running programs to share and collectively respond to the events generated by these input devices.

In this example, each input device (e.g., input devices A, B, BA, and BB, etc.) is managed by a software driver program hosted on the respective device (e.g., device A, device B, etc.)

which translates the discrete raw data generated by the input device hardware into Plasma proteins and deposits those proteins into a Plasma pool. For example, input device A generates raw data and provides the raw data to device A which, in turn, translates the discrete raw data into proteins (e.g., protein 1A, protein 2A, etc.) and deposits those proteins into the pool. As another example, input device BB generates raw data and provides the raw data to device B which, in turn, translates the discrete raw data into proteins (e.g., protein 1B, protein 3B, etc.) and deposits those proteins into the pool.

Each protein contains a descrip list that specifies the action registered by the input device as well as identifying information for the device itself. Where possible, the protein describes may also ascribe a general semantic meaning for the device action. The protein's data payload (e.g., ingests) carries the full set of useful state information for the device event.

To illustrate, here are example proteins for two typical events in such a system. Proteins are represented here as text however, in an actual implementation, the constituent parts of these proteins are typed data bundles (e.g., slaw). The protein describing a g-speak "one finger click" pose (described in the Related Applications) is as follows:

```
[ Descrips: { point, engage, one, one-finger-engage, hand,
  pilot-id-02, hand-id-23 }
  Ingests: { pilot-id => 02,
    hand-id => 23,
    pos => [ 0.0, 0.0, 0.0 ]
    angle-axis => [ 0.0, 0.0, 0.0, 0.707 ]
    gripe => .. ^:vx
    time => 184437103.29}]
```

As a further example, the protein describing a mouse click is as follows:

```
[ Descrips: { point, click, one, mouse-click, button-one,
  mouse-id-02 }
  Ingests: { mouse-id => 23,
    pos => [ 0.0, 0.0, 0.0 ]
    time => 184437124.80}]
```

Either or both of the sample proteins foregoing might cause a participating program of a host device to run a particular portion of its code. These programs may be interested in the general semantic labels: the most general of all, "point", or the more specific pair, "engage, one". Or they may be looking for events that would plausibly be generated only by a precise device: "one-finger-engage", or even a single aggregate object, "hand-id-23".

The proteins, as described above, are available in the pool for use by any program or device coupled or connected to the pool, regardless of type of the program or device. Consequently, any number of programs running on any number of computers may extract event proteins from the input pool. These devices need only be able to participate in the pool via either the local memory bus or a network connection in order to extract proteins from the pool. An immediate consequence of this is the beneficial possibility of decoupling processes that are responsible for generating 'input events' from those that use or interpret the events. Another consequence is the multiplexing of sources and consumers of events so that input devices may be controlled by one person or may be used simultaneously by several people (e.g., a Plasma-based input

framework supports many concurrent users), while the resulting event streams are in turn visible to multiple event consumers.

As an example of protein use, device C can extract one or more proteins (e.g., protein 1B, etc.) from the pool. Following protein extraction, device C can use the data of the protein, retrieved or read from the slaw of the descripts and ingests of the protein, in processing input events of input devices CA and CC to which the protein data corresponds. As another example, device A can extract one or more proteins (e.g., protein 1B, etc.) from the pool. Following protein extraction, device A can use the data of the protein in processing input events of input device A to which the protein data corresponds.

Devices and/or programs coupled or connected to a pool may skim backwards and forwards in the pool looking for particular sequences of proteins. It is often useful, for example, to set up a program to wait for the appearance of a protein matching a certain pattern, then skim backwards to determine whether this protein has appeared in conjunction with certain others. This facility for making use of the stored event history in the input pool often makes writing state management code unnecessary, or at least significantly reduces reliance on such undesirable coding patterns.

Examples of input devices that are used in the embodiments of the system described herein include gestural input sensors, keyboards, mice, infrared remote controls such as those used in consumer electronics, and task-oriented tangible media objects, to name a few.

FIG. 24 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow the numerous running programs to share and collectively respond to the graphics events generated by the devices, under yet another alternative embodiment. This system is but one example of a system comprising multiple running programs (e.g. graphics A-E) and one or more display devices (not shown), in which the graphical output of some or all of the programs is made available to other programs in a coordinated manner using the Plasma constructs (e.g., pools, proteins, and slaw) to allow the running programs to share and collectively respond to the graphics events generated by the devices.

It is often useful for a computer program to display graphics generated by another program. Several common examples include video conferencing applications, network-based slideshow and demo programs, and window managers. Under this configuration, the pool is used as a Plasma library to implement a generalized, framework which encapsulates video, network application sharing, and window management, and allows programmers to add in a number of features not commonly available in current versions of such programs.

Programs (e.g., graphics A-E) running in the Plasma compositing environment participate in a coordination pool through couplings and/or connections to the pool. Each program may deposit proteins in that pool to indicate the availability of graphical sources of various kinds. Programs that are available to display graphics also deposit proteins to indicate their displays' capabilities, security and user profiles, and physical and network locations.

Graphics data also may be transmitted through pools, or display programs may be pointed to network resources of other kinds (RTSP streams, for example). The phrase "graphics data" as used herein refers to a variety of different representations that lie along a broad continuum; examples of graphics data include but are not limited to literal examples

(e.g., an 'image', or block of pixels), procedural examples (e.g., a sequence of 'drawing' directives, such as those that flow down a typical OpenGL pipeline), and descriptive examples (e.g., instructions that combine other graphical constructs by way of geometric transformation, clipping, and compositing operations).

On a local machine graphics data may be delivered through platform-specific display driver optimizations. Even when graphics are not transmitted via pools, often a periodic screen-capture will be stored in the coordination pool so that clients without direct access to the more esoteric sources may still display fall-back graphics.

One advantage of the system described here is that unlike most message passing frameworks and network protocols, pools maintain a significant buffer of data. So programs can rewind backwards into a pool looking at access and usage patterns (in the case of the coordination pool) or extracting previous graphics frames (in the case of graphics pools).

FIG. 25 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow stateful inspection, visualization, and debugging of the running programs, under still another alternative embodiment. This system is but one example of a system comprising multiple running programs (e.g. program P-A, program P-B, etc.) on multiple devices (e.g., device A, device B, etc.) in which some programs access the internal state of other programs using or via pools.

Most interactive computer systems comprise many programs running alongside one another, either on a single machine or on multiple machines and interacting across a network. Multi-program systems can be difficult to configure, analyze and debug because run-time data is hidden inside each process and difficult to access. The generalized framework and Plasma constructs of an embodiment described herein allow running programs to make much of their data available via pools so that other programs may inspect their state. This framework enables debugging tools that are more flexible than conventional debuggers, sophisticated system maintenance tools, and visualization harnesses configured to allow human operators to analyze in detail the sequence of states that a program or programs has passed through.

Referring to FIG. 25, a program (e.g., program P-A, program P-B, etc.) running in this framework generates or creates a process pool upon program start up. This pool is registered in the system almanac, and security and access controls are applied. More particularly, each device (e.g., device A, B, etc.) translates discrete raw data generated by or output from the programs (e.g., program P-A, program P-B, etc.) running on that respective device into Plasma proteins and deposits those proteins into a Plasma pool. For example, program P-A generates data or output and provides the output to device A which, in turn, translates the raw data into proteins (e.g., protein 1A, protein 2A, protein 3A, etc.) and deposits those proteins into the pool. As another example, program P-B generates data and provides the data to device B which, in turn, translates the data into proteins (e.g., proteins 1B-4B, etc.) and deposits those proteins into the pool.

For the duration of the program's lifetime, other programs with sufficient access permissions may attach to the pool and read the proteins that the program deposits; this represents the basic inspection modality, and is a conceptually "one-way" or "read-only" proposition: entities interested in a program P-A inspect the flow of status information deposited by P-A in its process pool. For example, an inspection program or application running under device C can extract one or more pro-

teins (e.g., protein 1A, protein 2A, etc.) from the pool. Following protein extraction, device C can use the data of the protein, retrieved or read from the slaw of the descripts and ingests of the protein, to access, interpret and inspect the internal state of program P-A.

But, recalling that the Plasma system is not only an efficient stateful transmission scheme but also an omnidirectional messaging environment, several additional modes support program-to-program state inspection. An authorized inspection program may itself deposit proteins into program P's process pool to influence or control the characteristics of state information produced and placed in that process pool (which, after all, program P not only writes into but reads from).

FIG. 26 is a block diagram of a processing environment including multiple devices coupled among numerous programs running on one or more of the devices in which the Plasma constructs (e.g., pools, proteins, and slaw) are used to allow influence or control the characteristics of state information produced and placed in that process pool, under an additional alternative embodiment. In this system example, the inspection program of device C can for example request that programs (e.g., program P-A, program P-B, etc.) dump more state than normal into the pool, either for a single instant or for a particular duration. Or, prefiguring the next 'level' of debug communication, an interested program can request that programs (e.g., program P-A, program P-B, etc.) emit a protein listing the objects extant in its runtime environment that are individually capable of and available for interaction via the debug pool. Thus informed, the interested program can 'address' individuals among the objects in the programs runtime, placing proteins in the process pool that a particular object alone will take up and respond to. The interested program might, for example, request that an object emit a report protein describing the instantaneous values of all its component variables. Even more significantly, the interested program can, via other proteins, direct an object to change its behavior or its variables' values.

More specifically, in this example, inspection application of device C places into the pool a request (in the form of a protein) for an object list (e.g., "Request-Object List") that is then extracted by each device (e.g., device A, device B, etc.) coupled to the pool. In response to the request, each device (e.g., device A, device B, etc.) places into the pool a protein (e.g., protein 1A, protein 1B, etc.) listing the objects extant in its runtime environment that are individually capable of and available for interaction via the debug pool.

Thus informed via the listing from the devices, and in response to the listing of the objects, the inspection application of device C addresses individuals among the objects in the programs runtime, placing proteins in the process pool that a particular object alone will take up and respond to. The inspection application of device C can, for example, place a request protein (e.g., protein "Request Report P-A-O", "Request Report P-B-O") in the pool that an object (e.g., object P-A-O, object P-B-O, respectively) emit a report protein (e.g., protein 2A, protein 2B, etc.) describing the instantaneous values of all its component variables. Each object (e.g., object P-A-O, object P-B-O) extracts its request (e.g., protein "Request Report P-A-O", "Request Report P-B-O", respectively) and, in response, places a protein into the pool that includes the requested report (e.g., protein 2A, protein 2B, respectively). Device C then extracts the various report proteins (e.g., protein 2A, protein 2B, etc.) and takes subsequent processing action as appropriate to the contents of the reports.

In this way, use of Plasma as an interchange medium tends ultimately to erode the distinction between debugging, process control, and program-to-program communication and coordination.

To that last, the generalized Plasma framework allows visualization and analysis programs to be designed in a loosely-coupled fashion. A visualization tool that displays memory access patterns, for example, might be used in conjunction with any program that outputs its basic memory reads and writes to a pool. The programs undergoing analysis need not know of the existence or design of the visualization tool, and vice versa.

The use of pools in the manners described above does not unduly affect system performance. For example, embodiments have allowed for depositing of several hundred thousand proteins per second in a pool, so that enabling even relatively verbose data output does not noticeably inhibit the responsiveness or interactive character of most programs.

Embodiments described herein include a system comprising an input device. The system of an embodiment comprises a detector coupled to a processor and detecting an orientation of the input device. The input device of an embodiment has a plurality of modal orientations corresponding to the orientation. The plurality of modal orientations of an embodiment corresponds to a plurality of input modes of a gestural control system. The detector of an embodiment is coupled to the gestural control system and automatically controls selection of an input mode of the plurality of input modes in response to the orientation.

Embodiments described herein include a system comprising: an input device; and a detector coupled to a processor and detecting an orientation of the input device, wherein the input device has a plurality of modal orientations corresponding to the orientation, wherein the plurality of modal orientations correspond to a plurality of input modes of a gestural control system, wherein the detector is coupled to the gestural control system and automatically controls selection of an input mode of the plurality of input modes in response to the orientation.

The detector of an embodiment interprets and translates the orientation of the input device into input signals of the gestural control system.

The detector of an embodiment interprets and translates orientation transitions of the input device between the plurality of modal orientations into input signals of the gestural control system.

The detecting of the orientation of an embodiment comprises detecting an absolute three-space location of an instantaneous state of the input device at a point in time and space.

The detector of an embodiment tracks orientation transitions of the input device between the plurality of modal orientations.

The orientation transitions of an embodiment are rotational transitions about an axis of the input device.

The orientation transitions of an embodiment are rotational transitions about a plurality of axes of the input device.

The system of an embodiment comprises applying hysteresis to the selection of the input mode during the orientation transitions.

The detector of an embodiment tracks instantaneous orientation in real-time.

The detector of an embodiment translates raw tracking data of the input device into six degrees of spatial orientation.

The detector of an embodiment tracks instantaneous position of the input device in real-time.

The detector of an embodiment translates raw tracking data of the input device into six degrees of spatial position.

The system of an embodiment comprises at least one input sensor, wherein the input sensor is positioned on the input device.

The detector of an embodiment translates raw input sensor position data into input sensor state.

The detector of an embodiment translates raw input sensor position data into input sensor transition data.

Position of the input sensor of an embodiment controls the plurality of input modes.

The input device of an embodiment comprises the detector.

The detector of an embodiment is a remote detector that is remote to the input device.

The detector of an embodiment comprises an input device detector component carried by the input device and a remote detector component that is remote to the input device.

The input device of an embodiment is a hand-held input device.

The detector of an embodiment detects an absolute three-space location of an instantaneous state of the input device at a point in time and space.

The plurality of input modes of an embodiment comprise a direct manipulation mode in which the instantaneous state is used for direct manipulation of an application element of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprise a meta-manipulation mode in which the instantaneous state is used for direct manipulation of a set of application elements of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprise a three-dimensional manipulation mode in which the instantaneous state is used for three-dimensional manipulation of an application element of a component coupled to the gestural control system.

The gestural control system of an embodiment controls a three-space object coupled to the gestural control system.

The system of an embodiment comprises controlling the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

The controlling of an embodiment comprises a direct coupling between motion of the input device and the three-space object.

The controlling of an embodiment includes an indirect coupling between motion of the input device and the three-space object.

The three-space object of an embodiment is presented on a display device coupled to the processor.

The system of an embodiment comprises controlling movement of the three-space object by mapping the plurality of input modes of the input device to a plurality of object translations of the three-space object.

The mapping of an embodiment includes a direct mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes an indirect mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes correlating positional offsets of the plurality of input modes to positional offsets of the object translations of the three-space object.

The mapping of an embodiment includes correlating positional offsets of the input device to translational velocity of the object translations of the three-space object.

The system of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a linear translation of the three-space object.

The system of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a rotational translation of the three-space object.

The system of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a rotational translation of the three-space object.

The system of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a linear translation of the three-space object.

The detecting of an embodiment comprises detecting when an extrapolated position of the input device intersects virtual space, wherein the virtual space comprises space depicted on a display device coupled to the gestural control system.

The detector of an embodiment detects an event of the input device, wherein the event corresponds to at least one of a three-space location and a three-space orientation of the input device, wherein the event corresponds to an application of a first type.

The processor of an embodiment generates data sequences comprising input device event data specifying the event and state information of the event, wherein the input device event data and state information are type-specific data having a type corresponding to a first application of the gestural control system.

The processor of an embodiment forms a data capsule to include the data sequences, wherein the data capsule has a data structure comprising an application-independent representation of the data sequences.

The processor of an embodiment places the data capsule in a repository.

A second event of an embodiment running under an application of a second type searches the repository and identifies a correspondence between the data capsule and the second event. The application of the second type of an embodiment executes an operation corresponding to the second event using contents of the data sequences of the data capsule.

The tracking of an embodiment is electromagnetic field (EMF) tracking.

The input device of an embodiment includes the detector coupled to circuitry, wherein the detector comprises a plurality of coils.

The system of an embodiment comprises a field generator remote to the input device, wherein the field generator generates an EMF that induces signals in the plurality of coils.

The system of an embodiment comprises a plurality of field generators remote to the input device, wherein each field generator induces signals in the plurality of coils of the input device when the input device is proximate to the field generator that induces the signals.

The detector of an embodiment detects the orientation of the input device using the EMF signals induced in the plurality of coils.

The detector of an embodiment detects a position of the input device using the signals induced in the plurality of coils.

The input device of an embodiment comprises a transmitter coupled to the processor, wherein the transmitter is a wireless transmitter.

The transmitter of an embodiment communicates the orientation of the input device to the gestural control system.

The transmitter of an embodiment communicates a position of the input device to the gestural control system.

The input device of an embodiment comprises at least one input sensor, wherein the transmitter communicates state of the at least one input sensor to the gestural control system.

The tracking of an embodiment is optical tracking.

The system of an embodiment comprises at least one tag connected to the input device.

The at least one tag of an embodiment comprises a plurality of tags connected to a front region of the input device.

The tracking of an embodiment includes dynamically detecting a position of the at least one tag.

The tracking of an embodiment includes detecting position of a set of tags coupled to a region of the input device.

Each tag of the set of tags of an embodiment includes a pattern, wherein each pattern of each tag of the set of tags is different than any pattern of any remaining tag of the plurality of tags.

Each tag of an embodiment includes a first pattern and a second pattern, wherein the first pattern is common to any tag of the set of tags and the second pattern is different between at least two tags of the set of tags.

The set of tags of an embodiment form a plurality of patterns on the input device.

The at least one tag of an embodiment comprises a set of infrared (IR) light-emitting diodes (LEDs) and a set of retro-reflective dots.

The input device of an embodiment includes a microprocessor coupled to circuitry, wherein the circuitry is coupled to the set of IR LEDs.

The system of an embodiment comprises at least one input sensor coupled to the circuitry, wherein the at least one input sensor is positioned on the input device.

The at least one input sensor of an embodiment controls a state of the set of IR LEDs.

The state of each IR LED of an embodiment corresponds to at least one input mode of the plurality of input modes.

The at least one tag of an embodiment comprises at least one tracking dot connected to the input device.

The at least one tracking dot of an embodiment comprises at least one infrared (IR) light-emitting diode (LED).

The at least one tracking dot of an embodiment comprises at least one retro-reflective dot.

The at least one tracking dot of an embodiment comprises at least one of an infrared (IR) light-emitting diode (LED) and a retro-reflective dot.

The at least one tracking dot of an embodiment comprises a set of IR LEDs and a set of retro-reflective dots.

The input device of an embodiment comprises a plurality of sides, wherein each side of the plurality of sides corresponds to a modal orientation of the plurality of modal orientations.

Each side of the plurality of sides of an embodiment corresponds to an input mode of the plurality of input modes.

Each side of the plurality of sides of an embodiment is assigned an input mode to which the side corresponds.

Each side of the plurality of sides of an embodiment is dynamically assigned an input mode based on a context.

At least one side of the plurality of sides of an embodiment comprises an input sensor.

Each side of the plurality of sides of an embodiment comprises an input sensor.

Each input sensor of each side of an embodiment is dynamically assigned a function based on the orientation.

Each input sensor of each side of an embodiment is dynamically assigned a function based on a context.

The plurality of sides of an embodiment includes three sides, wherein the input device has a triangular cross-section, wherein each side of the plurality of sides corresponds to a modal orientation of the plurality of modal orientations.

Each side of the plurality of sides of an embodiment corresponds to an input mode of the plurality of input modes.

A first side of the plurality of sides of an embodiment corresponds to a first input mode and a second side of the plurality of sides of an embodiment corresponds to a second input mode. An orientation transition of an embodiment comprises rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

The detector of an embodiment applies a hysteresis band to the selection of the input mode during the orientation transition.

The hysteresis band of an embodiment is at least one of equal to and greater than approximately 30 degrees.

The hysteresis band of an embodiment is programmable.

Each side of the plurality of sides of an embodiment comprises an input sensor.

Embodiments described herein include a system comprising an input device including a processor coupled to circuitry.

The processor of an embodiment controls a plurality of input modes of a gestural control system. The system of an embodiment comprises a detector coupled to the processor and detecting at least one of an absolute three-space position and orientation of an instantaneous state of the input device at a point in time and space. The input device of an embodiment has a plurality of modal orientations corresponding to at least the orientation. The plurality of modal orientations of an embodiment corresponds to a plurality of input modes of a gestural control system. The detector of an embodiment is coupled to the gestural control system and automatically controls selection of an input mode of the plurality of input modes in response to the orientation.

Embodiments described herein include a system comprising: an input device comprising a processor coupled to circuitry, wherein the processor controls a plurality of input modes of a gestural control system; and a detector coupled to the processor and detecting at least one of an absolute three-space position and orientation of an instantaneous state of the input device at a point in time and space, wherein the input device has a plurality of modal orientations corresponding to at least the orientation, wherein the plurality of modal orientations correspond to a plurality of input modes of a gestural control system, wherein the detector is coupled to the gestural control system and automatically controls selection of an input mode of the plurality of input modes in response to the orientation.

Embodiments described herein include an input device comprising a hand-held housing that includes a processor. The processor of an embodiment is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device. The input signals of an embodiment automatically control a plurality of input modes of a gestural control system.

Embodiments described herein include an input device comprising a hand-held housing that includes a processor, wherein the processor is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device, wherein the input signals automatically control a plurality of input modes of a gestural control system.

The detector of an embodiment comprises circuitry including a plurality of coils.

The device of an embodiment comprises a remote electromagnetic field (EMF) generator that generates an EMF that induces signals in the plurality of coils.

The device of an embodiment comprises a plurality of remote EMF generators, wherein each EMF generator induces signals in the plurality of coils when the housing is proximate to the EMF generator that induces the signals.

The detector of an embodiment detects the orientation using the EMF signals induced in the plurality of coils.

The detector of an embodiment detects a position of the input device using the signals induced in the plurality of coils.

The device of an embodiment comprises a transmitter coupled to the processor, wherein the transmitter is a wireless transmitter.

The transmitter of an embodiment communicates the orientation to the gestural control system.

The transmitter of an embodiment communicates a position of the input device to the gestural control system.

The device of an embodiment comprises at least one input sensor coupled to the processor, wherein the transmitter communicates state of the at least one input sensor to the gestural control system.

The detector of an embodiment translates raw input sensor position data into input sensor state.

Position of the input sensor of an embodiment controls the gestural control system.

The plurality of orientations of an embodiment corresponds to the plurality of input modes of a gestural control system.

The plurality of input modes of an embodiment comprise a direct manipulation mode in which the input device directly manipulates of an application element of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprises a meta-manipulation mode in which the input device directly manipulates a set of application elements of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprises a three-dimensional manipulation mode in which the input device controls three-dimensional manipulation of an application element of a component coupled to the gestural control system.

The input signals of an embodiment automatically control selection of an input mode of the plurality of input modes in response to the orientation.

The device of an embodiment comprises applying hysteresis to selection of the input mode during the transitions.

The transitions of an embodiment are rotational transitions about at least one axis of the input device.

The detector of an embodiment tracks in real-time instantaneous three-space orientation at a point in time and space.

The detector of an embodiment translates raw tracking data of the input device into six degrees of spatial orientation.

The detector of an embodiment tracks in real-time instantaneous three-space position at a point in time and space.

The detector of an embodiment translates raw tracking data of the input device into six degrees of spatial position.

The housing of an embodiment comprises a plurality of sides, wherein each side of the plurality of sides corresponds to an orientation of the plurality of orientations.

Each side of the plurality of sides of an embodiment corresponds to an input mode of the plurality of input modes.

Each side of the plurality of sides of an embodiment is assigned an input mode to which the side corresponds.

Each side of the plurality of sides of an embodiment is dynamically assigned an input mode based on a context.

At least one side of the plurality of sides of an embodiment comprises an input sensor.

Each side of the plurality of sides of an embodiment comprises an input sensor.

Each input sensor of each side of an embodiment is dynamically assigned a function based on the orientation.

Each input sensor of each side of an embodiment is dynamically assigned a function based on a context.

The plurality of sides of an embodiment includes three sides, wherein the housing has a triangular cross-section, wherein each side of the plurality of sides corresponds to an orientation of the plurality of orientations.

A first side of the plurality of sides of an embodiment corresponds to a first input mode and a second side of the plurality of sides corresponds to a second input mode, wherein a transition comprising rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

The detector of an embodiment applies a hysteresis band to the selection of the input mode during the transition.

The hysteresis band of an embodiment is at least one of equal to and greater than approximately 30 degrees.

The hysteresis band of an embodiment is programmable.

The detector of an embodiment detects an event of the input device, wherein the event corresponds to at least one of a three-space location and a three-space orientation of the input device, wherein the event corresponds to an application of a first type.

The processor of an embodiment generates data sequences comprising input device event data specifying the event and state information of the event, wherein the input device event data and state information are type-specific data having a type corresponding to a first application of the gestural control system.

The processor of an embodiment forms a data capsule to include the data sequences, wherein the data capsule has a data structure comprising an application-independent representation of the data sequences.

The processor of an embodiment places the data capsule in a repository.

A second event running under an application of a second type of an embodiment searches the repository and identifies a correspondence between the data capsule and the second event, wherein the application of the second type executes an operation corresponding to the second event using contents of the data sequences of the data capsule.

The gestural control system of an embodiment controls a three-space object coupled to the gestural control system.

The device of an embodiment comprises controlling via the device the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

The controlling of an embodiment comprises a direct coupling between motion of the input device and the three-space object.

The controlling of an embodiment includes an indirect coupling between motion of the input device and the three-space object.

The three-space object of an embodiment is presented on a display device coupled to the processor.

The device of an embodiment comprises controlling movement of the three-space object by mapping the plurality of input modes of the input device to a plurality of object translations of the three-space object.

The mapping of an embodiment includes a direct mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes an indirect mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes correlating positional offsets of the plurality of input modes to positional offsets of the object translations of the three-space object.

The mapping of an embodiment includes correlating positional offsets of the input device to translational velocity of the object translations of the three-space object.

The device of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a linear translation of the three-space object.

The device of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a rotational translation of the three-space object.

The device of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a rotational translation of the three-space object.

The device of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a linear translation of the three-space object.

The detecting of an embodiment comprises detecting when an extrapolated position of the input device intersects virtual space, wherein the virtual space comprises space depicted on a display device coupled to the gestural control system.

Embodiments described herein include an input device comprising a processor coupled to a detector and contained in a housing having a hand-held form factor. The detector of an embodiment detects and translates into input signals at least one of an orientation in which the housing is currently operated and a position of the housing. The detector of an embodiment detects the at least one of the orientation and the position using signals induced in coils. The input device of an embodiment comprises a transmitter coupled to the processor. The transmitter of an embodiment communicates the input signals to a gestural control system. The input signals of an embodiment automatically control a plurality of input modes of the gestural control system and control translation and rotation of a three-space object coupled to the gestural control system.

Embodiments described herein include an input device comprising: a processor coupled to a detector and contained in a housing having a hand-held form factor, wherein the detector detects and translates into input signals at least one of an orientation in which the housing is currently operated and a position of the housing, wherein the detector detects the at least one of the orientation and the position using signals induced in coils; and a transmitter coupled to the processor, wherein the transmitter communicates the input signals to a gestural control system, wherein the input signals automatically control a plurality of input modes of the gestural control system and control translation and rotation of a three-space object coupled to the gestural control system.

Embodiments described herein include an input device comprising a hand-held housing that includes a processor. The processor of an embodiment is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device. The input signals of an embodiment automatically control a plurality of input modes of a gestural control system and control a three-space object coupled to the gestural control system. The input signals of an embodiment control the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

Embodiments described herein include an input device comprising a hand-held housing that includes a processor,

wherein the processor is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device, wherein the input signals automatically control a plurality of input modes of a gestural control system and control a three-space object coupled to the gestural control system, wherein the input signals control the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

Embodiments described herein include a method comprising detecting and translating an orientation in which an input device is currently operated and transitions between a plurality of orientations of the input device. The input device of an embodiment comprises a hand-held housing that includes a processor. The method of an embodiment comprises automatically controlling selection of an input mode of the plurality of input modes in response to the orientation and the transitions. The method of an embodiment comprises controlling a three-space object coupled to the gestural control system according to the input mode.

Embodiments described herein include a method comprising: detecting and translating an orientation in which an input device is currently operated and transitions between a plurality of orientations of the input device, wherein the input device comprises a hand-held housing that includes a processor; automatically controlling selection of an input mode of the plurality of input modes in response to the orientation and the transitions; and controlling a three-space object coupled to the gestural control system according to the input mode.

The detecting of the orientation of an embodiment comprises detecting an absolute three-space location of an instantaneous state of the input device at a point in time and space.

The transitions of an embodiment are rotational transitions about an axis of the input device.

The transitions of an embodiment are rotational transitions about a plurality of axes of the input device.

The method of an embodiment comprises applying hysteresis to the selection of the input mode during the transitions.

The detecting of an embodiment comprises detecting instantaneous orientation in real-time.

The translating of an embodiment comprises translating raw tracking data of the input device into six degrees of spatial orientation.

The detecting of an embodiment comprises detecting instantaneous position of the input device in real-time.

The translating of an embodiment comprises translating raw tracking data of the input device into six degrees of spatial position.

The method of an embodiment comprises at least one input sensor, wherein the input sensor is positioned on the input device.

The method of an embodiment comprises detecting and translating into input sensor state raw input sensor position data of an input sensor of the input device.

The method of an embodiment comprises controlling the plurality of input modes in response to the input sensor state.

The detecting of an embodiment is performed onboard the input device.

The detecting of an embodiment is performed remote to the input device.

The detecting of an embodiment is performed onboard the input device and performed remote to the input device.

The input device of an embodiment is a hand-held input device.

The detecting of an embodiment comprises detecting an absolute three-space location of an instantaneous state of the input device at a point in time and space.

The plurality of input modes of an embodiment comprise a direct manipulation mode in which the instantaneous state is used for direct manipulation of an application element of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprise a meta-manipulation mode in which the instantaneous state is used for direct manipulation of a set of application elements of a component coupled to the gestural control system.

The plurality of input modes of an embodiment comprise a three-dimensional manipulation mode in which the instantaneous state is used for three-dimensional manipulation of an application element of a component coupled to the gestural control system.

The method of an embodiment comprises controlling via the gestural control system a three-space object coupled to the gestural control system.

The method of an embodiment comprises controlling the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

The controlling of an embodiment comprises a direct coupling between motion of the input device and the three-space object.

The controlling of an embodiment includes an indirect coupling between motion of the input device and the three-space object.

The method of an embodiment comprises presenting the three-space object on a display device.

The method of an embodiment comprises controlling movement of the three-space object by mapping the plurality of input modes of the input device to a plurality of object translations of the three-space object.

The mapping of an embodiment includes a direct mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes an indirect mapping between the plurality of input modes and the plurality of object translations.

The mapping of an embodiment includes correlating positional offsets of the plurality of input modes to positional offsets of the object translations of the three-space object.

The mapping of an embodiment includes correlating positional offsets of the input device to translational velocity of the object translations of the three-space object.

The method of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a linear translation of the three-space object.

The method of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a rotational translation of the three-space object.

The method of an embodiment comprises controlling movement of the three-space object by mapping a linear gesture of the input device to a rotational translation of the three-space object.

The method of an embodiment comprises controlling movement of the three-space object by mapping a rotational gesture of the input device to a linear translation of the three-space object.

The detecting of an embodiment comprises detecting when an extrapolated position of the input device intersects virtual space, wherein the virtual space comprises space depicted on a display device coupled to the gestural control system.

The detecting of an embodiment comprises an event of the input device, wherein the event corresponds to at least one of a three-space location and a three-space orientation of the input device, wherein the event corresponds to an application of a first type.

The method of an embodiment comprises generating data sequences comprising input device event data specifying the event and state information of the event, wherein the input device event data and state information are type-specific data having a type corresponding to a first application of the gestural control system.

The method of an embodiment comprises forming a data capsule to include the data sequences, wherein the data capsule has a data structure comprising an application-independent representation of the data sequences.

The method of an embodiment comprises placing the data capsule in a repository.

A second event running under an application of a second type of an embodiment searches the repository and identifies a correspondence between the data capsule and the second event, wherein the application of the second type executes an operation corresponding to the second event using contents of the data sequences of the data capsule.

The detecting of an embodiment uses electromagnetic field (EMF) tracking.

The method of an embodiment comprises generating an EMF that induces signals in a plurality of coils of the input device.

The method of an embodiment comprises generating an EMF using a plurality of field generators remote to the input device, wherein each field generator induces signals in the plurality of coils of the input device when the input device is proximate to the field generator that induces the signals.

The detecting of an embodiment comprises detecting the orientation of the input device using the EMF signals induced in the plurality of coils.

The detecting of an embodiment comprises detecting a position of the input device using the signals induced in the plurality of coils.

The method of an embodiment comprises communicating the orientation of the input device to the gestural control system.

The method of an embodiment comprises communicating a position of the input device to the gestural control system.

The method of an embodiment comprises communicating a state of at least one input sensor of the input device to the gestural control system.

The detecting of an embodiment uses optical tracking.

The method of an embodiment comprises tracking at least one tag connected to the input device.

The at least one tag of an embodiment comprises a plurality of tags connected to a front region of the input device.

The tracking of an embodiment includes dynamically detecting a position of the at least one tag.

The tracking of an embodiment includes detecting position of a set of tags coupled to a region of the input device.

Each tag of the set of tags of an embodiment includes a pattern, wherein each pattern of each tag of the set of tags is different than any pattern of any remaining tag of the plurality of tags.

Each tag of an embodiment includes a first pattern and a second pattern, wherein the first pattern is common to any tag of the set of tags and the second pattern is different between at least two tags of the set of tags.

The set of tags of an embodiment form a plurality of patterns on the input device.

The at least one tag of an embodiment comprises a set of infrared (IR) light-emitting diodes (LEDs) and a set of retro-reflective dots.

The method of an embodiment comprises controlling a state of the set of IR LEDs via at least one input sensor of the input device.

The state of each IR LED of an embodiment corresponds to at least one input mode of the plurality of input modes.

The at least one tag of an embodiment comprises at least one tracking dot connected to the input device.

The at least one tracking dot of an embodiment comprises at least one infrared (IR) light-emitting diode (LED).

The at least one tracking dot of an embodiment comprises at least one retro-reflective dot.

The at least one tracking dot of an embodiment comprises at least one of an infrared (IR) light-emitting diode (LED) and a retro-reflective dot.

The at least one tracking dot of an embodiment comprises a set of IR LEDs and a set of retro-reflective dots.

The method of an embodiment comprises associating each side of a plurality of sides of the input device with an orientation of a plurality of orientations.

The method of an embodiment comprises defining a correspondence between each side of the plurality of sides and an input mode of the plurality of input modes.

The method of an embodiment comprises assigning each side of the plurality of sides to an input mode.

The method of an embodiment comprises dynamically assigning each side of the plurality of sides to an input mode based on a context.

The method of an embodiment comprises dynamically assigning a function to an input sensor of a side based on the orientation.

The method of an embodiment comprises dynamically assigning a function to an input sensor of each side based on a context.

The plurality of sides of an embodiment includes three sides, wherein the input device has a triangular cross-section, comprising associating each side of the plurality of sides to an orientation of the plurality of orientations.

The method of an embodiment comprises defining a correspondence between each side of the plurality of sides and an input mode of the plurality of input modes.

The correspondence of an embodiment comprises a first side of the plurality of sides corresponding to a first input mode and a second side of the plurality of sides corresponding to a second input mode, wherein the transition comprising rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

The method of an embodiment comprises applying a hysteresis band to the selection of the input mode during the transition.

The hysteresis band of an embodiment is at least one of equal to and greater than approximately 30 degrees.

The hysteresis band of an embodiment is programmable.

Embodiments described herein include a method comprising automatically detecting a gesture of an input device from gesture data received via a detector. The gesture data of an embodiment is absolute three-space orientation data of an instantaneous state of the input device at a point in time and space. The method of an embodiment comprises identifying the gesture using only the gesture data. The method of an embodiment comprises translating the gesture to a gesture signal. The method of an embodiment comprises automati-

cally controlling selection of an input mode of a plurality of input modes of the input device in response to the gesture signal.

Embodiments described herein include a method comprising: automatically detecting a gesture of an input device from gesture data received via a detector, wherein the gesture data is absolute three-space orientation data of an instantaneous state of the input device at a point in time and space, and identifying the gesture using only the gesture data; translating the gesture to a gesture signal; and automatically controlling selection of an input mode of a plurality of input modes of the input device in response to the gesture signal.

Embodiments described herein include a method comprising detecting an orientation of an input device having a plurality of modal orientations corresponding to the orientation. The plurality of modal orientations of an embodiment corresponds to a plurality of input modes of a gestural control system. The method of an embodiment comprises automatically controlling selection of an input mode of the plurality of input modes in response to the orientation. The method of an embodiment comprises controlling a three-space object coupled to the gestural control system according to the input mode.

Embodiments described herein include a method comprising: detecting an orientation of an input device having a plurality of modal orientations corresponding to the orientation, wherein the plurality of modal orientations correspond to a plurality of input modes of a gestural control system; automatically controlling selection of an input mode of the plurality of input modes in response to the orientation; and controlling a three-space object coupled to the gestural control system according to the input mode.

The systems and methods described herein include and/or run under and/or in association with a processing system. The processing system includes any collection of processor-based devices or computing devices operating together, or components of processing systems or devices, as is known in the art. For example, the processing system can include one or more of a portable computer, portable communication device operating in a communication network, and/or a network server. The portable computer can be any of a number and/or combination of devices selected from among personal computers, cellular telephones, personal digital assistants, portable computing devices, and portable communication devices, but is not so limited. The processing system can include components within a larger computer system.

The processing system of an embodiment includes at least one processor and at least one memory device or subsystem. The processing system can also include or be coupled to at least one database. The term "processor" as generally used herein refers to any logic processing unit, such as one or more central processing units (CPUs), digital signal processors (DSPs), application-specific integrated circuits (ASIC), etc. The processor and memory can be monolithically integrated onto a single chip, distributed among a number of chips or components of a host system, and/or provided by some combination of algorithms. The methods described herein can be implemented in one or more of software algorithm(s), programs, firmware, hardware, components, circuitry, in any combination.

System components embodying the systems and methods described herein can be located together or in separate locations. Consequently, system components embodying the systems and methods described herein can be components of a single system, multiple systems, and/or geographically separate systems. These components can also be subcomponents or subsystems of a single system, multiple systems, and/or

geographically separate systems. These components can be coupled to one or more other components of a host system or a system coupled to the host system.

Communication paths couple the system components and include any medium for communicating or transferring files among the components. The communication paths include wireless connections, wired connections, and hybrid wireless/wired connections. The communication paths also include couplings or connections to networks including local area networks (LANs), metropolitan area networks (MANs), wide area networks (WANs), proprietary networks, interoffice or backend networks, and the Internet. Furthermore, the communication paths include removable fixed mediums like floppy disks, hard disk drives, and CD-ROM disks, as well as flash RAM, Universal Serial Bus (USB) connections, RS-232 connections, telephone lines, buses, and electronic mail messages.

Unless the context clearly requires otherwise, throughout the description, the words "comprise," "comprising," and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of "including, but not limited to." Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words "herein," "hereunder," "above," "below," and words of similar import refer to this application as a whole and not to any particular portions of this application. When the word "or" is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

The above description of embodiments of the processing environment is not intended to be exhaustive or to limit the systems and methods described to the precise form disclosed. While specific embodiments of, and examples for, the processing environment are described herein for illustrative purposes, various equivalent modifications are possible within the scope of other systems and methods, as those skilled in the relevant art will recognize. The teachings of the processing environment provided herein can be applied to other processing systems and methods, not only for the systems and methods described above.

The elements and acts of the various embodiments described above can be combined to provide further embodiments. These and other changes can be made to the processing environment in light of the above detailed description.

In general, in the following claims, the terms used should not be construed to limit the embodiments to the specific embodiments disclosed in the specification and the claims, but should be construed to include all systems that operate under the claims. Accordingly, the embodiments are not limited by the disclosure herein, but instead the scope of the embodiments is to be determined entirely by the claims.

While certain aspects of the embodiments are presented below in certain claim forms, the inventors contemplate the various aspects of the embodiments in any number of claim forms. Accordingly, the inventors reserve the right to add additional claims after filing the application to pursue such additional claim forms for other aspects of the embodiments.

What is claimed is:

1. An input device comprising: a processor coupled to a detector and contained in a housing having a hand-held form factor, wherein the detector detects and translates into input signals at least one of an orientation in which the housing is currently operated and a position of the housing, wherein the detector detects the at least one of the orientation and the position using signals induced in coils; and a transmitter

coupled to the processor, wherein the transmitter communicates the input signals to a gestural control system, wherein the input signals automatically control a plurality of input modes of the gestural control system and control translation and rotation of a three-space object coupled to the gestural control system; wherein a first side of a plurality of sides corresponds to a first input mode and a second side of a plurality of sides corresponds to a second input mode, wherein a transition comprising rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

2. An input device comprising a hand-held housing that includes a processor, wherein the processor is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device, wherein the input signals automatically control a plurality of input modes of a gestural control system and control a three-space object coupled to the gestural control system, wherein the input signals control the three-space object through three translational degrees of freedom and three rotational degrees of freedom; wherein a first side of a plurality of sides corresponds to a first input mode and a second side of a plurality of sides corresponds to a second input mode, wherein a transition comprising rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

3. An input device comprising a hand-held housing that includes a processor, wherein the processor is coupled to a detector that detects and translates into input signals an orientation in which the input device is currently operated and transitions between a plurality of orientations of the input device, wherein the input signals automatically control a plurality of input modes of a gestural control system; wherein a first side of a plurality of sides corresponds to a first input mode and a second side of a plurality of sides corresponds to a second input mode, wherein a transition comprising rotating the input device around a longitudinal axis more than 120 degrees relative to a center of the first side causes a change from the first input mode to the second input mode.

4. The device of claim 3, wherein the detector comprises circuitry including a plurality of coils.

5. The device of claim 4, comprising a remote electromagnetic field (EMF) generator that generates an EMF that induces signals in the plurality of coils.

6. The device of claim 5, comprising a plurality of remote EMF generators, wherein each EMF generator induces signals in the plurality of coils when the housing is proximate to the EMF generator that induces the signals.

7. The device of claim 5, wherein the detector detects the orientation using the EMF signals induced in the plurality of coils.

8. The device of claim 7, wherein the detector detects a position of the input device using the signals induced in the plurality of coils.

9. The device of claim 4, comprising a transmitter coupled to the processor, wherein the transmitter is a wireless transmitter.

10. The device of claim 9, wherein the transmitter communicates the orientation to the gestural control system.

11. The device of claim 9, wherein the transmitter communicates a position of the input device to the gestural control system.

61

12. The device of claim 9, comprising at least one input sensor coupled to the processor, wherein the transmitter communicates state of the at least one input sensor to the gestural control system.

13. The device of claim 12, wherein the detector translates raw input sensor position data into input sensor state.

14. The device of claim 13, wherein position of the input sensor controls the gestural control system.

15. The device of claim 3, wherein the plurality of orientations corresponds to the plurality of input modes of a gestural control system.

16. The device of claim 15, wherein the plurality of input modes comprise a direct manipulation mode in which the input device directly manipulates of an application element of a component coupled to the gestural control system.

17. The device of claim 15, wherein the plurality of input modes comprise a meta-manipulation mode in which the input device directly manipulates a set of application elements of a component coupled to the gestural control system.

18. The device of claim 15, wherein the plurality of input modes comprise a three-dimensional manipulation mode in which the input device controls three-dimensional manipulation of an application element of a component coupled to the gestural control system.

19. The device of claim 3, wherein the input signals automatically control selection of an input mode of the plurality of input modes in response to the orientation.

20. The device of claim 19, comprising applying hysteresis to selection of the input mode during the transitions.

21. The device of claim 3, wherein the transitions are rotational transitions about at least one axis of the input device.

22. The device of claim 3, wherein the detector tracks in real-time instantaneous three-space orientation at a point in time and space.

23. The device of claim 22, wherein the detector translates raw tracking data of the input device into six degrees of spatial orientation.

24. The device of claim 3, wherein the detector tracks in real-time instantaneous three-space position at a point in time and space.

25. The device of claim 24, wherein the detector translates raw tracking data of the input device into six degrees of spatial position.

26. The device of claim 3, wherein the housing comprises a plurality of sides, wherein each side of the plurality of sides corresponds to an orientation of the plurality of orientations.

27. The device of claim 26, wherein each side of the plurality of sides corresponds to an input mode of the plurality of input modes.

28. The device of claim 27, wherein each side of the plurality of sides is assigned an input mode to which the side corresponds.

29. The device of claim 27, wherein each side of the plurality of sides is dynamically assigned an input mode based on a context.

30. The device of claim 26, wherein at least one side of the plurality of sides comprises an input sensor.

31. The device of claim 26, wherein each side of the plurality of sides comprises an input sensor.

32. The device of claim 31, wherein each input sensor of each side is dynamically assigned a function based on the orientation.

33. The device of claim 31, wherein each input sensor of each side is dynamically assigned a function based on a context.

62

34. The device of claim 26, wherein the plurality of sides includes three sides, wherein the housing has a triangular cross-section, wherein each side of the plurality of sides corresponds to an orientation of the plurality of orientations.

35. The device of claim 20, wherein the detector applies a hysteresis band to the selection of the input mode during the transition.

36. The device of claim 35, wherein the hysteresis band is at least one of equal to and greater than approximately 30 degrees.

37. The device of claim 35, wherein the hysteresis band is programmable.

38. The device of claim 3, wherein the detector detects an event of the input device, wherein the event corresponds to at least one of a three-space location and a three-space orientation of the input device, wherein the event corresponds to an application of a first type.

39. The device of claim 38, wherein the processor generates data sequences comprising input device event data specifying the event and state information of the event, wherein the input device event data and state information are type-specific data having a type corresponding to a first application of the gestural control system.

40. The device of claim 39, wherein the processor forms a data capsule to include the data sequences, wherein the data capsule has a data structure comprising an application-independent representation of the data sequences.

41. The device of claim 40, wherein the processor places the data capsule in a repository.

42. The device of claim 41, wherein a second event running under an application of a second type searches the repository and identifies a correspondence between the data capsule and the second event, wherein the application of the second type executes an operation corresponding to the second event using contents of the data sequences of the data capsule.

43. The device of claim 3, wherein the gestural control system controls a three-space object coupled to the gestural control system.

44. The device of claim 43, comprising controlling via the device the three-space object through three translational degrees of freedom and three rotational degrees of freedom.

45. The device of claim 44, wherein the controlling comprises a direct coupling between motion of the input device and the three-space object.

46. The device of claim 44, wherein the controlling includes an indirect coupling between motion of the input device and the three-space object.

47. The device of claim 44, wherein the three-space object is presented on a display device coupled to the processor.

48. The device of claim 44, comprising controlling movement of the three-space object by mapping the plurality of input modes of the input device to a plurality of object translations of the three-space object.

49. The device of claim 48, wherein the mapping includes a direct mapping between the plurality of input modes and the plurality of object translations.

50. The device of claim 48, wherein the mapping includes an indirect mapping between the plurality of input modes and the plurality of object translations.

51. The device of claim 48, wherein the mapping includes correlating positional offsets of the plurality of input modes to positional offsets of the object translations of the three-space object.

52. The device of claim 48, wherein the mapping includes correlating positional offsets of the input device to translational velocity of the object translations of the three-space object.

53. The device of claim **44**, comprising controlling movement of the three-space object by mapping a linear gesture of the input device to a linear translation of the three-space object.

54. The device of claim **44**, comprising controlling movement of the three-space object by mapping a rotational gesture of the input device to a rotational translation of the three-space object. 5

55. The device of claim **44**, comprising controlling movement of the three-space object by mapping a linear gesture of the input device to a rotational translation of the three-space object. 10

56. The device of claim **44**, comprising controlling movement of the three-space object by mapping a rotational gesture of the input device to a linear translation of the three-space object. 15

57. The device of claim **44**, wherein the detecting comprises detecting when an extrapolated position of the input device intersects virtual space, wherein the virtual space comprises space depicted on a display device coupled to the gestural control system. 20

* * * * *